



**Juniper Networks
EX2500 Ethernet Switch**

Configuration Guide

Release 3.0

Juniper Networks, Inc.

1194 North Mathilda Avenue
Sunnyvale, CA 94089

USA

408-745-2000

www.juniper.net

Juniper Networks, the Juniper Networks logo, JUNOS, NetScreen, ScreenOS, and Steel-Belted Radius are registered trademarks of Juniper Networks, Inc. in the United States and other countries. JUNOS is a trademark of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

Juniper Networks EX2500 Ethernet Switch Configuration Guide, Release 3.0
Copyright © 2009 Juniper Networks, Inc.
All rights reserved.

Writing: William Rogers
Editing: Taffy Everts
Illustration: William Rogers

Revision History
26 October 2009—Revision 2
11 April 2009—Revision 1

The information in this document is current as of the date listed in the revision history.

SOFTWARE LICENSE

The terms and conditions for using this software are described in the software license contained in the acknowledgment to your purchase order or, to the extent applicable, to any reseller agreement or end-user purchase agreement executed between you and Juniper Networks. By using this software, you indicate that you understand and agree to be bound by those terms and conditions.

Generally speaking, the software license restricts the manner in which you are permitted to use the software and may contain prohibitions against certain uses. The software license may state conditions under which the license is automatically terminated. You should consult the license for further details.

For complete product documentation, see the Juniper Networks Web site at <http://www.juniper.net/techpubs>.

Table of Contents

About This Guide	xi
Objectives	xi
Audience	xi
Supported Platforms	xi
Documentation Conventions.....	xii
List of Technical Publications	xiii
Documentation Feedback	xiii
Requesting Technical Support.....	xiii
Self-Help Online Tools and Resources.....	xiv
Opening a Case with JTAC	xiv

Part 1

EX2500 Ethernet Switch Applications

Chapter 1	Accessing the Switch	3
	Configuring the Management Interface	3
	Dynamic Host Configuration Protocol	4
	Using Telnet.....	5
	Using the EX2500 Web Device Manager	5
	Configuring EX2500 Web Device Manager Access via HTTP	6
	Configuring EX2500 Web Device Manager Access via HTTPS.....	6
	Using SNMP	7
	SNMPv1, SNMPv2.....	7
	SNMPv3.....	8
	Default Configuration.....	8
	User Configuration	8
	Configuring SNMP Trap Hosts.....	9
	SNMPv1 Trap Host Configuration.....	9
	SNMPv2 Trap Host Configuration.....	9
	SNMPv3 Trap Host Configuration.....	10
	Securing Access to the Switch	10
	RADIUS Authentication and Authorization.....	11
	How RADIUS Authentication Works.....	11
	Configuring RADIUS on the Switch.....	11
	RADIUS Authentication Features in the EX2500 Switch	12
	Switch User Accounts.....	13
	RADIUS Attributes for EX2500 User Privileges	13

	TACACS + Authentication.....	14
	How TACACS + Authentication Works.....	14
	TACACS + Authentication Features in the EX2500 Switch	14
	Command Authorization and Logging.....	16
	Configuring TACACS + Authentication on the Switch.....	16
	Secure Shell	17
	Configuring SSH Features on the Switch	17
	SSH Encryption of Management Messages.....	17
	Generating RSA Host and Server Keys for SSH Access	17
	SSH Integration with RADIUS and TACACS + Authentication	18
	End User Access Control.....	18
	Considerations for Configuring End User Accounts	19
	User Access Control	19
	Listing Current Users.....	20
	Logging In to an End User Account	20
Chapter 2	VLANs	21
	VLAN Overview.....	21
	VLANs and Port VLAN ID Numbers	22
	VLAN Numbers.....	22
	PVID Numbers.....	22
	VLAN Tagging	23
	VLAN Topologies and Design Considerations	26
	VLAN Configuration Rules	26
	Multiple VLANs Configuration Example	27
	Private VLANs	29
	Private VLAN Ports	29
	Private VLAN Configuration Guidelines.....	30
	Private VLAN Configuration Example	30
Chapter 3	Spanning Tree Protocol	31
	Spanning Tree Overview	31
	Bridge Protocol Data Units (BPDUs).....	32
	Determining the Path for Forwarding BPDUs.....	32
	Bridge Priority.....	32
	Port Priority	33
	Port Path Cost.....	33
	Spanning Tree Group Configuration Guidelines	33
	Changing the Spanning Tree Mode.....	33
	Assigning a VLAN to a Spanning Tree Group	33
	Creating a VLAN.....	34
	Rules for VLAN Tagged Ports.....	34
	Adding and Removing Ports from STGs	34
	Rapid Spanning Tree Protocol.....	35
	Port State Changes	35
	Port Type and Link Type	36
	Edge Port	36
	Link Type.....	36
	RSTP Configuration Guidelines	36
	RSTP Configuration Example.....	36
	Per VLAN Rapid Spanning Tree.....	37
	Default Spanning Tree Configuration.....	37
	Why Do We Need Multiple Spanning Trees?.....	37

	PVRST Configuration Guidelines	38
	Configuring PVRST	38
	Multiple Spanning Tree Protocol	39
	MSTP Region	39
	Common Internal Spanning Tree.....	39
	MSTP Configuration Guidelines.....	39
	Multiple Spanning Tree Groups Configuration Example	40
	Fast Uplink Convergence	41
	Configuration Guidelines	42
	Configuring Fast Uplink Convergence	42
Chapter 4	Ports and Trunking	43
	Trunking Overview	43
	Statistical Load Distribution	43
	Built-In Fault Tolerance.....	44
	Before Configuring Static Trunks	44
	Trunk Group Configuration Rules.....	44
	Port Trunking Configuration Example	45
	Configurable Trunk Hash Algorithm	47
	Link Aggregation Control Protocol.....	47
	LACP Configuration Guidelines	49
	Configuring LACP	49
	Optionally Reducing LACP Timeout	49
Chapter 5	Quality of Service	51
	QoS Overview	51
	Using ACL Filters.....	52
	MAC Extended ACLs	53
	IP Standard ACLs	53
	IP Extended ACLs	54
	Understanding ACL Priority	55
	Assigning ACLs to a Port.....	56
	Viewing ACL Statistics	56
	ACL Configuration Examples	56
	ACL Example 1—Blocking Traffic to a Host.....	56
	ACL Example 2—Blocking Traffic from a Source to a Destination	57
	ACL Example 3—Blocking HTTP Traffic	57
	ACL Example 4—Blocking All Except Certain Packets	58
	Using Storm Control Filters	59
	Broadcast Storms	59
	Configuring Storm Control.....	59
	Using DSCP Values to Provide QoS	60
	Differentiated Services Concepts	60
	Per Hop Behavior	61
	QoS Levels.....	62
	DSCP Mapping.....	62
	Using 802.1p Priority to Provide QoS.....	63
	Queuing and Scheduling	64
Chapter 6	Remote Monitoring	65
	RMON Overview	65
	RMON Group 1—Statistics	66
	RMON Group 2—History	67

	History MIB Object ID.....	67
	Configuring RMON History	67
	RMON Group 3—Alarms	68
	Alarm MIB Objects.....	68
	Configuring RMON Alarms	68
	RMON Group 9—Events.....	69
Chapter 7	IGMP	71
	IGMP Snooping	71
	FastLeave	72
	IGMPv3 Snooping.....	73
	IGMP Snooping Configuration Example	73
	Static Multicast Router	74
Chapter 8	High Availability Through Uplink Failure Detection	75
	High Availability Overview	75
	Failure Detection Pair.....	76
	Spanning Tree Protocol with UFD	76
	UFD Configuration Guidelines	76
	UFD Configuration Example.....	77
	Monitoring UFD	77
Part 2	Appendixes	
Appendix A	Monitoring Ports with Port Mirroring	81
	Port Mirroring Overview	81
	Configuring Port Mirroring	82
Part 3	Indexes	
	Index.....	85

List of Figures

Figure 1: Default VLAN Settings.....	24
Figure 2: Port-Based VLAN Assignment	25
Figure 3: 802.1Q Tagging (after Port-Based VLAN Assignment)	25
Figure 4: 802.1Q Tag Assignment	25
Figure 5: 802.1Q Tagging (after 802.1Q Tag Assignment)	26
Figure 6: Sample Network with Multiple VLANs.....	27
Figure 7: Two VLANs on One Spanning Tree Group.....	37
Figure 8: Two VLANs, Each on a Different Spanning Tree Group	38
Figure 9: Implementing Multiple Spanning Tree Groups	40
Figure 10: Port Trunk Group Configuration Example	45
Figure 11: QoS Model	52
Figure 12: Layer 3 IPv4 Packet	60
Figure 13: Layer 2 802.1q/802.1p VLAN-Tagged Packet	63
Figure 14: Uplink Failure Detection Example.....	76
Figure 15: Monitoring Ports	81

List of Tables

Table 1:	Notice Icons.....	xii
Table 2:	EX2500 Text and Syntax Conventions.....	xii
Table 3:	EX2500 Ethernet Switch Documentation.....	xiii
Table 4:	User Access Levels.....	13
Table 5:	EX2500-Proprietary Attributes for RADIUS.....	13
Table 6:	Default TACACS+ Authorization Levels.....	15
Table 7:	Alternate TACACS+ Authorization Levels.....	15
Table 8:	Components of Sample Network with Multiple VLANs.....	27
Table 9:	Ports, Trunk Groups, and VLANs.....	32
Table 10:	RSTP vs. STP Port States.....	35
Table 11:	Actor vs. Partner LACP Configuration.....	48
Table 12:	Well-Known Protocol Types.....	54
Table 13:	Well-Known Application Ports.....	55
Table 14:	Assured Forwarding.....	61
Table 15:	Class Selector.....	61
Table 16:	Default QoS Service Levels.....	62

About This Guide

This preface provides the following guidelines for using the *Juniper Networks EX2500 Ethernet Switch Configuration Guide*:

- Objectives on page xi
- Audience on page xi
- Supported Platforms on page xi
- Documentation Conventions on page xii
- List of Technical Publications on page xiii
- Documentation Feedback on page xiii
- Requesting Technical Support on page xiii

Objectives

This guide describes how to configure and use the software on the EX2500 Ethernet Switch.



NOTE: This guide documents Release 3.0 of the EX2500 Ethernet Switch. For additional information—either corrections to or information that might have been omitted from this manual—see the *EX2500 Ethernet Switch 3.0 Release Notes* at <http://www.juniper.net/>.

Audience

This guide is intended for network installers and system administrators engaged in configuring and maintaining a network. Administrators must be familiar with Ethernet concepts, IP addressing, the IEEE 802.1D Spanning Tree Protocol (STP), and SNMP configuration.

Supported Platforms

The features described in this guide are supported by only the EX2500 software running on EX2500 Ethernet Switches only.

Documentation Conventions

Table 1 describes the notice icons used in this manual.

Table 2 describes the EX2500 text and syntax conventions.

Table 1: Notice Icons





Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.

Table 2: EX2500 Text and Syntax Conventions

Convention	Usage	Examples
Bold text like this	Names of windows, dialog boxes, buttons, tabs, and other objects in a user interface that you click or select.	Click the Submit button on the bottom of the form.
Bold text like this	In syntax descriptions and set-off command examples, text you must type exactly as shown	Main# sys
Braces { }	Required elements in syntax that has more than one option. You must choose one of the options. Do not type the braces.	show portchannel {<1-12> hash information} (For example, you can enter either show portchannel 3, show portchannel hash, or show portchannel information.)
Brackets []	Optional elements in syntax descriptions. Do not type the brackets.	copy running config tftp [data-port mgt-port] (You enter either copy running config tftp data-port, copy running config tftp mgt-port, or copy running config tftp.)
Fixed-width text like this	Onscreen computer output	ex2500(config)# reload Reset will use software "image2"...
<i>Italic text like this</i>	Book titles, special terms, and words to be emphasized	See the <i>EX2500 Ethernet Switch Command Reference</i> .
<Italic text like this in angle brackets>	Variables in command syntax. Replace the italic text with the appropriate real name or value when entering the command. Do not type the brackets.	To establish a Telnet session, enter host# telnet <IP address> (For example, you can enter telnet 192.32.10.12 .)
Plain text like this	Names of commands, files, and directories used within the text	View the readme.txt file.
Vertical line	Separates choices for command keywords and arguments. Enter only one choice. Do not type the vertical line.	copy running config tftp [data-port mgt-port] (You enter either copy running config tftp data-port, copy running config tftp mgt-port, or copy running config tftp.)

List of Technical Publications

Table 3 lists the documentation supporting the EX2500 Ethernet Switch. All documentation for EX Series Ethernet Switches is available at <http://www.juniper.net/techpubs/>.

Table 3: EX2500 Ethernet Switch Documentation

Document	Description
<i>EX2500 Ethernet Switch Quick Start</i>	Provides brief installation and initial configuration instructions.
<i>EX2500 Ethernet Switch Hardware Guide</i>	Provides information and instructions for installing an EX2500 Ethernet Switch.
<i>EX2500 Ethernet Switch Web Device Manager Guide</i>	Provides an overview of how to access and use the EX2500 Web Device Manager.
<i>EX2500 Ethernet Switch Configuration Guide</i>	Describes how to configure and use the software on the EX2500 Ethernet Switch.
<i>EX2500 Ethernet Switch Command Reference</i>	Describes how to configure and use the software with your EX2500 Ethernet Switch. The reference lists each command and includes the complete syntax and a functional description, using the EX2500 command-line interface (CLI)
<i>EX2500 Ethernet Switch 3.0 Release Notes</i>	Summarize EX2500 switch features and known problems, provide information that might have been omitted from the manuals, and provide upgrade and downgrade instructions.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. Send e-mail to techpubs-comments@juniper.net with the following information:

- Document URL or title
- Page number
- Software version
- Your name and company

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract, or are covered under warranty, and need postsales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <http://www.juniper.net/customers/support/downloads/710059.pdf>.
- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, the Juniper Networks online self-service portal—the Customer Support Center (CSC)—provides the following features:

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes: <http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <http://www.juniper.net/alerts/>
- Join and participate in the Juniper Networks Community Forum: <http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

To verify service entitlement by product and serial number, use our Serial Number Entitlement (SNE) Tool at <http://tools.juniper.net/SerialNumber/EntitlementSearch/>.

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <http://www.juniper.net/cm/>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, visit us at <http://www.juniper.net/support/requesting-support.html>.

Part 1

EX2500 Ethernet Switch Applications

This configuration guide will help you plan, implement, and administer EX2500 software. Where possible, each chapter provides feature overviews, usage examples, and configuration instructions.

- “Accessing the Switch” on page 3 describes how to access the switch to perform administration tasks. This chapter also discusses different methods to manage the switch for remote administrators using specific IP addresses, authentication, and Secure Shell (SSH).
- “VLANs” on page 21 describes how to configure Virtual Local Area Networks (VLANs) for creating separate network segments, including how to use VLAN tagging for devices that use multiple VLANs. This chapter also describes private VLANs.
- “Spanning Tree Protocol” on page 31 discusses how Spanning Trees configure the network so that the switch uses the most efficient path when multiple paths exist.
- “Ports and Trunking” on page 43 describes how to group multiple physical ports together to aggregate the bandwidth between large-scale network devices.
- “Quality of Service” on page 51 discusses Quality of Service features, including IP filtering using Access Control Lists, Differentiated Services, and IEEE 802.1p priority values.
- “Remote Monitoring” on page 65 discusses how to configure and use the Remote Monitoring (RMON) agent on the switch.
- “IGMP” on page 71,” describes how the EX2500 software implements IGMP Snooping to handle multicast traffic efficiently.
- “High Availability Through Uplink Failure Detection” on page 75 describes how to use Uplink Failure Detection (UFD) to ensure that network resources remain available if one switch is removed for service.

Chapter 1

Accessing the Switch

The EX2500 software provides a means for accessing, configuring, and viewing information and statistics about the EX2500 Ethernet Switch. This chapter discusses different methods of accessing the switch and ways to secure the switch for remote administrators:

- Configuring the Management Interface on page 3
- Dynamic Host Configuration Protocol on page 4
- Using Telnet on page 5
- Using the EX2500 Web Device Manager on page 5
- Using SNMP on page 7
- Securing Access to the Switch on page 10

Configuring the Management Interface

To manage the switch through the management port, you must configure an IP interface. Configure the following IP parameters:

- IP address
- Subnet mask
- Default gateway address

To configure the IP parameters, use the following procedure:

1. Log on to the switch.
2. Enter Global Configuration mode.

```
ex2500> enable
ex2500# configure terminal
```

3. Configure the management IP address, subnet mask, and default gateway.

```
ex2500(config)# interface ip-mgmt address 10.10.10.2
ex2500(config)# interface ip-mgmt netmask 255.255.255.0
ex2500(config)# interface ip-mgmt enable
ex2500(config)# interface ip-mgmt gateway 10.10.10.1
ex2500(config)# interface ip-mgmt gateway enable
ex2500(config)# exit
```

Once you configure the IP address for your switch, you can connect to the management port and use the Telnet program from an external management station to access and control the switch. The management port provides *out-of-band* management.

You also can configure *in-band* management through any of the switch data ports. To allow in-band management, use the following procedure:

1. Log on to the switch.
2. Enter IP interface mode.

```
ex2500> enable
ex2500# configure terminal
ex2500(config)# interface ip 1
```

3. Configure the management IP interface, subnet mask, and VLAN assignment. Enable the interface.

```
ex2500(config-ip-if)# ip address 10.10.10.2
ex2500(config-ip-if)# ip netmask 255.255.255.0
ex2500(config-ip-if)# ipvlan 1
ex2500(config-ip-if)# enable
```

4. Configure the default gateway. Enable the gateway.

```
ex2500(config-ip-if)# ip gateway address 10.10.10.1
ex2500(config-ip-if)# ip gateway enable
ex2500(config-ip-if)# exit
```

Once you configure the IP address and you have an existing network connection, you can use the Telnet program from an external management station to access and control the switch. Once the default gateway is enabled, the management station and your switch do not need to be on the same IP subnet.

Dynamic Host Configuration Protocol

Dynamic Host Configuration Protocol (DHCP) is a transport protocol that provides a framework for automatically assigning IP addresses and configuration information to other IP hosts or clients in a large TCP/IP network. Without DHCP, the IP address must be entered manually for each network device. DHCP allows a network administrator to distribute IP addresses from a central point and automatically send a new IP address when a device is connected to a different place in the network.

DHCP is an extension of another network IP management protocol, Bootstrap Protocol (BOOTP), with an additional capability of being able to allocate reusable network addresses and configuration parameters for client operation.

Built on the client/server model, DHCP allows hosts or clients on an IP network to obtain their configurations from a DHCP server, thereby reducing network administration. The most significant configuration the client receives from the server is its required IP address; (other optional parameters include the “generic” filename to be booted, the address of the default gateway, and so forth).

To enable DHCP on an IP interface, use the following commands:

```
ex2500(config)# interface ip 1
ex2500(config-ip-if)# dhcp enable
ex2500(config-ip-if)# exit
```

Using Telnet

A Telnet connection offers the convenience of accessing the switch from any workstation connected to the network. Telnet access provides the same options for user access and administrator access as those available through the console port.

To configure the switch for Telnet access, you need to have a device with Telnet software located on the same network as the switch. The switch must have an IP address. The switch can get its IP address in one of two ways:

- Dynamically, from a DHCP server on your network
- Manually, when you configure the switch IP address

Once you have configured the switch with an IP address and gateway, you can access the switch from any workstation connected to the management network. Telnet access provides the same options for user and administrator access as those available through the console port.

By default, Telnet access is enabled. Use the following command to disable or enable Telnet access:

```
ex2500(config)# [no] access telnet
```

To establish a Telnet connection with the switch, you can run the Telnet program on your workstation and issue the **telnet** command with the switch IP address:

```
telnet <switch IP address>
```

Using the EX2500 Web Device Manager

The EX2500 Web Device Manager is a Web-based management interface for interactive switch access through your Web browser.

The Web Device Manager provides access to the common configuration, management and operation features of the switch through your Web browser. For more information, see the *EX2500 Ethernet Switch Web Device Manager Guide*.

By default, EX2500 Web Device Manager access is enabled on the switch.

Configuring EX2500 Web Device Manager Access via HTTP

By default, EX2500 Web Device Manager access via HTTP is **enabled**. Use the following command to disable or enable EX2500 Web Device Manager access on the switch via HTTP:

```
ex2500(config)# [no] access http enable
```

The default HTTP Web server port to access the EX2500 Web Device Manager is port 80. However, you can change the default Web server port with the following command:

```
ex2500(config)# access http port <TCP port number>
```

For workstation access to your switch via the EX2500 Web Device Manager, open a Web browser window and type in the URL using the IP interface address of the switch, such as:

```
http://10.10.10.1
```

Configuring EX2500 Web Device Manager Access via HTTPS

The EX2500 Web Device Manager can be accessed via a secure HTTPS connection over management and data ports. By default, EX2500 Web Device Manager access via HTTPS is **disabled**.

To enable EX2500 Web Device Manager access on the switch via HTTPS, use the following command:

```
ex2500(config)# access https enable
```

To change the HTTPS Web server port number from the default port 443, use the following command:

```
ex2500(config)# access https port <TCP port number>
```

Accessing the EX2500 Web Device Manager via HTTPS requires a SSL certificate to be used during the key exchange. A default certificate is created the first time HTTPS is enabled, but you can import a new certificate that defines the information you want to be used. Use the following command to import the SSL certificate:

```
ex2500(config)# access https import-certificate
```

The certificate is saved to Flash memory for use once the switch is rebooted.

When a client (e.g. Web browser) connects to the switch, the client is asked to accept the certificate and verify that the fields match what is expected. Once EX2500 Web Device Manager access is granted to the client, the Web Device Manager can be used as described in the *EX2500 Ethernet Switch Web Device Manager Guide*.

The EX2500 Web Device Manager is organized at a high level as follows:

Context tabs—These tabs allow you to select the type of action you wish to perform. The *Configure* tab provides access to the configuration elements for the entire switch. The *Monitor* tab provides access to the switch statistics and state information. The *Dashboard* tab allows you to display settings and operating status of a variety of switch features.

Navigation Window—This window provides a menu list of switch features and functions, as follows:

- **System**—This folder provides access to the configuration elements for the entire switch.
- **Switch Ports**—Configure each of the physical ports on the switch.
- **Port-Based Port Mirroring**—Configure port mirroring and mirror ports.
- **Layer 2 Management**—Configure Layer 2 features, such as VLANs and Spanning Tree.
- **RMON Menu**—Configure Remote Monitoring (RMON).
- **Layer 3 Management**—Configure the switch interface, default gateway, and Internet Group Multicast Protocol (IGMP).
- **QoS**—Configure Quality of Service (QoS) features for the switch.
- **Access Control**—Configure Access Control Lists (ACLs) to filter IP packets.
- **Uplink Failure Detection**—Configure Uplink Failure Detection to provide high availability.

Using SNMP

The EX2500 switch provides SNMPv1, SNMPv2, and SNMPv3 support for access through any network management software, such as IBM Director or HP-OpenView.

SNMPv1, SNMPv2

To access the SNMP agent on the EX2500 switch, the read and write community strings on the SNMP manager should be configured to match those on the switch. The default read community string on the switch is **public**, and the default write community string is **private**.

The read and write community strings on the switch can be changed with the following commands on the CLI:

```
ex2500(config)# snmp-server read-community <1-32 characters>
```

and

```
ex2500(config)# snmp-server write-community <1-32 characters>
```

The SNMP manager should be able to reach the management interface or any one of the IP interfaces on the switch.

SNMPv3

SNMPv3 is an enhanced version of the Simple Network Management Protocol, approved by the Internet Engineering Steering Group in March, 2002. SNMPv3 contains additional security and authentication features that provide data origin authentication, data integrity checks, timeliness indicators, and encryption to protect against threats such as masquerade, modification of information, message stream modification, and disclosure.

SNMPv3 ensures that the client can use SNMPv3 to query the MIBs, mainly for security.

For more information on SNMP MIBs and the commands used to configure SNMP on the switch, see the *EX2500 Ethernet Switch Command Reference*.

Default Configuration

The EX2500 switch has two SNMPv3 users by default. Both of the following users have access to all the MIBs supported by the switch:

1. username 1: **adminmd5** (password adminmd5). Authentication used is MD5.
2. username 2: **adminsha** (password adminsha). Authentication used is SHA.

To configure an SNMP username, enter the following command from the CLI:

```
ex2500(config)# snmp-server user <1-16> name <1-32>
```

User Configuration

Users can be configured to use the authentication and privacy options. The EX2500 switch supports two authentication algorithms: MD5 and SHA, as specified in the following command:

```
ex2500(config)# snmp-server user <1-16> authentication-protocol md5 | sha
```

1. To configure a user with name admin, authentication type MD5, authentication password admin, and privacy option DES with a privacy password of admin, use the following CLI commands:

```
ex2500(config)# snmp-server user 5 name admin
ex2500(config)# snmp-server user 5 authentication-protocol md5
authentication-password
```

Changing authentication password; validation required:

Enter current admin password: <admin. password>

Enter new authentication password: <auth. password>

Re-enter new authentication password: <auth. password>

New authentication password accepted.

```
ex2500(config)# snmp-server user 5 privacy-protocol des privacy-password
```

Changing privacy password; validation required:

Enter current admin password: <admin. password>

Enter new privacy password: <privacy password>

Re-enter new privacy password: <privacy password>

New privacy password accepted.

2. Configure a user access group, along with the views the group may access. Use the access table to configure the group's access level. Because the read view, write view, and notify view are all set to iso, the user type has access to all private and public MIBs.

```
ex2500(config)# snmp-server access 5 name admingrp
ex2500(config)# snmp-server access 5 level authpriv
ex2500(config)# snmp-server access 5 read-view iso
ex2500(config)# snmp-server access 5 write-view iso
ex2500(config)# snmp-server access 5 notify-view iso
```

3. Assign the user to the user group. Use the group table to link the user to a particular access group.

```
ex2500(config)# snmp-server group 5 user-name admin
ex2500(config)# snmp-server group 5 group-name admingrp
```

Configuring SNMP Trap Hosts

SNMPv1 Trap Host Configuration

1. Configure an entry in the notify table.

```
ex2500(config)# snmp-server notify 10 name public
ex2500(config)# snmp-server notify 10 tag v1trap
```

2. Specify the IP address and other trap parameters in the targetAddr and targetParam tables. Use the following command to specify the username used with this targetParam table:

```
snmp-server target-parameters <1-16> user-name
```

For example:

```
ex2500(config)# snmp-server target-address 10 name v1trap address
10.70.70.190
ex2500(config)# snmp-server target-address 10 parameters-name v1param
ex2500(config)# snmp-server target-address 10 taglist v1param
ex2500(config)# snmp-server target-parameters 10 name v1param
ex2500(config)# snmp-server target-parameters 10 user-name v1only
ex2500(config)# snmp-server target-parameters 10 message snmpv1
```

SNMPv2 Trap Host Configuration

The SNMPv2 trap host configuration is similar to the SNMPv1 trap host configuration. Wherever you specify the model, use snmpv2 instead of snmpv1.

```
ex2500(config)# snmp-server read-community public
ex2500(config)# snmp-server target-address 1 name v2trap2 address 10.70.70.190
ex2500(config)# snmp-server target-address 1 parameters-name v2param2
ex2500(config)# snmp-server target-address 1 taglist v2param2
ex2500(config)# snmp-server target-parameters 1 name v2param2
ex2500(config)# snmp-server target-parameters 1 user-name v2only
ex2500(config)# snmp-server target-parameters 1 message snmpv2
ex2500(config)# snmp-server notify 1 name public
ex2500(config)# snmp-server notify 1 tag v2param2
```

SNMPv3 Trap Host Configuration

To configure a user for SNMPv3 traps, you can choose to send the traps with both privacy and authentication, with authentication only, or without privacy or authentication.

This is configured in the access table with the following commands:

```
ex2500(config)# snmp-server access <1-32> level
ex2500(config)# snmp-server target-parameters <1-16>
```

Configure the user in the user table accordingly.

It is not necessary to configure the community table for SNMPv3 traps because the community string is not used by SNMPv3.

The following example shows how to configure a SNMPv3 user v3trap with authentication only:

```
ex2500(config)# snmp-server user 11 name v3trap
ex2500(config)# snmp-server user 11 authentication-protocol md5
authentication-password
```

Changing authentication password; validation required:

Enter current admin password: *<admin. password>*

Enter new authentication password: *<auth. password>*

Re-enter new authentication password: *<auth. password>*

New authentication password accepted.

```
ex2500(config)# snmp-server access 11 notify-view iso
ex2500(config)# snmp-server access 11 level authnopriv
ex2500(config)# snmp-server group 11 user-name v3trap
ex2500(config)# snmp-server group 11 tag v3trap
ex2500(config)# snmp-server notify 11 name v3trap
ex2500(config)# snmp-server notify 11 tag v3trap
ex2500(config)# snmp-server target-address 11 name v3trap address 47.81.25.66
ex2500(config)# snmp-server target-address 11 taglist v3trap
ex2500(config)# snmp-server target-address 11 parameters-name v3param
ex2500(config)# snmp-server target-parameters 11 name v3param
ex2500(config)# snmp-server target-parameters 11 user-name v3trap
ex2500(config)# snmp-server target-parameters 11 level authNoPriv
```

Securing Access to the Switch

Secure switch management is needed for environments that perform significant management functions across the Internet. Common functions for secured management are described in the following sections:

- RADIUS Authentication and Authorization on page 11
- TACACS + Authentication on page 14
- End User Access Control on page 18

RADIUS Authentication and Authorization

The EX2500 switch supports the RADIUS (Remote Authentication Dial-in User Service) method to authenticate and authorize remote administrators for managing the switch. This method is based on a client/server model. The Remote Access Server (RAS)—the switch—is a client to the back-end database server. A remote user (the remote administrator) interacts only with the RAS, not the back-end server and database.

RADIUS authentication consists of the following components:

- A protocol with a frame format that utilizes UDP over IP (based on RFC 2138 and RFC 2866)
- A centralized server that stores all the user authorization information
- A client, in this case, the switch

The EX2500 switch—acting as the RADIUS client—communicates to the RADIUS server to authenticate and authorize a remote administrator using the protocol definitions specified in RFCs 2138 and 2866. Transactions between the client and the RADIUS server are authenticated by a shared key that is not sent over the network. In addition, the remote administrator passwords are sent encrypted between the RADIUS client (the switch) and the back-end RADIUS server.

How RADIUS Authentication Works

1. The remote administrator connects to the switch and provides username and password.
2. Using Authentication/Authorization protocol, the switch sends request to authentication server.
3. The authentication server checks the request against the user ID database.
4. Using RADIUS protocol, the authentication server instructs the switch to grant or deny administrative access.

Configuring RADIUS on the Switch

Use the following procedure to configure RADIUS authentication on your switch:

1. Configure the Primary and Secondary RADIUS servers, and enable RADIUS authentication.

```
ex2500(config)# radius-server primary-host 10.10.1.1
ex2500(config)# radius-server secondary-host 10.10.1.2
ex2500(config)# radius-server enable
```

2. Configure the RADIUS secret.

```
ex2500(config)# radius-server primary-host 10.10.1.1 key <1-32 character
secret>
ex2500(config)# radius-server secondary-host 10.10.1.2 key <1-32 character
secret>
```

3. If desired, you may change the default UDP port number used to listen to RADIUS. The well-known port for RADIUS is 1812.

```
ex2500(config)# radius-server port <UDP port number>
```

4. Configure the number of retry attempts for contacting the RADIUS server, and the timeout period.

```
ex2500(config)# radius-server retransmit 3
ex2500(config)# radius-server timeout 5
```

RADIUS Authentication Features in the EX2500 Switch

The EX2500 switch supports the following RADIUS authentication features:

- Supports a RADIUS client on the switch, based on the protocol definitions in RFC 2138 and RFC 2866.
- Allows a RADIUS secret password of up to 32 bytes and less than 16 octets.
- Supports a *secondary authentication server* so that when the primary authentication server is unreachable, the switch can send client authentication requests to the secondary authentication server. Use the following command to show the currently active RADIUS authentication server:

```
ex2500# show radius-server
```

- Supports user-configurable RADIUS server retry and time-out values:
 - Time-out value = 1 to 10 seconds
 - Retries = 1 through 3

The switch will time out if it does not receive a response from the RADIUS server in 1 through 3 retries. The switch will also automatically retry connecting to the RADIUS server before it declares the server down.

- Supports a user-configurable RADIUS application port. The default is 1812/UDP, based on RFC 2138. Port 1645 is also supported.
- Allows the network administrator to define privileges for one or more specific users to access the switch at the RADIUS user database.

Switch User Accounts

The user accounts listed in Table 4 can be defined in the RADIUS server dictionary file.

Table 4: User Access Levels

User Account	Description and Tasks Performed	Password
User	The User has no direct responsibility for switch management. He or she can view all switch status information and statistics but cannot make any configuration changes to the switch.	user
Operator	The Operator manages all functions of the switch. The Operator can reset ports, except the management port.	oper
Administrator	The super-user Administrator has complete access to all commands, information, and configuration commands on the switch, including the ability to change both the user and administrator passwords.	admin

RADIUS Attributes for EX2500 User Privileges

When the user logs in, the switch authenticates his or her level of access by sending the RADIUS access request, that is, the client authentication request, to the RADIUS authentication server.

If the remote user is successfully authenticated by the authentication server, the switch will verify the *privileges* of the remote user and authorize the appropriate access. The administrator has an option to allow *secure backdoor* access via Telnet, SSH, or the Web Device Manager. Secure backdoor provides switch access when the RADIUS servers cannot be reached. You always can access the switch via the console port, by using **noradius** and the administrator password, whether secure backdoor is enabled or not.



NOTE: To obtain the RADIUS backdoor password for your EX2500 switch, contact technical support.

All user privileges, other than those assigned to the Administrator, have to be defined in the RADIUS dictionary. RADIUS attribute 6, which is built into all RADIUS servers, defines the administrator. The filename of the dictionary is RADIUS vendor-dependent. Table 5 shows the RADIUS attributes defined for EX2500 user privilege levels.

Table 5: EX2500-Proprietary Attributes for RADIUS

Username/Access	User-Service-Type	Value
User	<i>Vendor-supplied</i>	255
Operator	<i>Vendor-supplied</i>	252
Admin	<i>Vendor-supplied</i>	6

TACACS+ Authentication

The EX2500 switch supports authentication and authorization with networks using the TACACS+ protocol. The EX2500 switch functions as the Network Access Server (NAS) by interacting with the remote client and initiating authentication and authorization sessions with the TACACS+ access server. The remote user is defined as someone requiring management access to the EX2500 switch either through a data port or a management port.

TACACS+ offers the following advantages over RADIUS:

- TACACS+ uses TCP-based connection-oriented transport, whereas RADIUS is UDP-based. TCP offers a connection-oriented transport, while UDP offers best-effort delivery. RADIUS requires additional programmable variables such as re-transmit attempts and time-outs to compensate for best-effort transport, but it lacks the level of built-in support that a TCP transport offers.
- TACACS+ offers full packet encryption, whereas RADIUS offers password-only encryption in authentication requests.
- TACACS+ separates authentication, authorization, and accounting.

How TACACS+ Authentication Works

TACACS+ works in much the same way as RADIUS authentication, as described on page 11. The remote administrator connects to the switch and provides a username and password.

1. Using Authentication/Authorization protocol, the switch sends a request to authentication server.
2. The authentication server checks the request against the user ID database.
3. Using TACACS+ protocol, the authentication server instructs the switch to grant or deny administrative access.

During a session, if additional authorization checking is needed, the switch checks with a TACACS+ server to determine if the user is granted permission to use a particular command.

TACACS+ Authentication Features in the EX2500 Switch

Authentication is the action of determining the identity of a user, and is generally done when the user first attempts to log in to a device or gain access to its services. The EX2500 switch supports ASCII inbound login to the device. PAP, CHAP, and ARAP login methods; TACACS+ change password requests; and one-time password authentication are not supported.

Authorization

Authorization is the action of determining a user's privileges on the device, and usually takes place after authentication.

The default mapping between TACACS+ authorization levels and EX2500 management access levels is shown in Table 6. The authorization levels must be defined on the TACACS+ server.

Table 6: Default TACACS+ Authorization Levels

EX2500 User Access Level	TACACS+ level
user	0
oper	3
admin	6

Alternate mapping between TACACS+ authorization levels and EX2500 management access levels is shown in Table 7. Use the following command to set the alternate TACACS+ authorization levels:

```
ex2500(config)# tacacs-server privilege-mapping
```

Table 7: Alternate TACACS+ Authorization Levels

EX2500 User Access Level	TACACS+ level
user	0 - 1
oper	6 - 8
admin	14 - 15

If the remote user is successfully authenticated by the authentication server, the switch verifies the *privileges* of the remote user and authorizes the appropriate access. The administrator has an option to allow *secure backdoor* access via Telnet or SSH. Secure backdoor provides switch access when the TACACS+ servers cannot be reached. You always can access the switch via the console port by using **notacacs** and the administrator password, whether secure backdoor is enabled or not.



NOTE: To obtain the TACACS+ backdoor password for your EX2500 switch, contact technical support.

Accounting

Accounting is the action of recording a user's activities on the device for the purposes of billing and/or security. It follows the authentication and authorization actions. If the authentication and authorization are not performed via TACACS+, no TACACS+ accounting messages are sent out. The EX2500 switch supports the following TACACS+ accounting attributes:

- protocol (console, telnet, ssh, or http)
- start_time
- stop_time
- elapsed_time
- disc_cause



NOTE: When you are using the EX2500 Web Device Manager, the TACACS+ Accounting Stop records are sent only if the **Logout** button on the browser is clicked.

Command Authorization and Logging

When TACACS+ Command Authorization is enabled, EX2500 configuration commands are sent to the TACACS+ server for authorization. Use the following command to enable TACACS+ Command Authorization:

```
ex2500(config)# tacacs-server command-authorization
```

When TACACS+ Command Logging is enabled, EX2500 configuration commands are logged on the TACACS+ server. Use the following command to enable TACACS+ Command Logging:

```
ex2500(config)# tacacs-server command-logging
```

The following examples illustrate the format of EX2500 commands sent to the TACACS+ server:

```
authorization request, cmd=shell, cmd-arg=interface ip
accounting request, cmd=shell, cmd-arg=interface ip
authorization request, cmd=shell, cmd-arg=enable
accounting request, cmd=shell, cmd-arg=enable
```

Configuring TACACS+ Authentication on the Switch

1. Configure the Primary and Secondary TACACS+ servers, and enable TACACS authentication.

```
ex2500(config)# tacacs-server primary-host 10.10.1.1
ex2500(config)# tacacs-server secondary-host 10.10.1.2
ex2500(config)# tacacs-server enable
```

2. Configure the TACACS+ secret and second secret.

```
ex2500(config)# tacacs-server primary-host 10.10.1.1 key <1-32 character
secret>
ex2500(config)# tacacs-server secondary-host 10.10.1.2 key <1-32 character
secret>
```

3. If desired, you may change the default TCP port number used to listen to TACACS+. The well-known port for TACACS+ is 49.

```
ex2500(config)# tacacs-server port <TCP port number>
```

4. Configure the number of retry attempts and the timeout period.

```
ex2500(config)# tacacs-server retransmit 3
ex2500(config)# tacacs-server timeout 5
```

Secure Shell

Secure Shell (SSH) uses secure tunnels to encrypt and secure messages between a remote administrator and the switch. Telnet does not provide this level of security. The Telnet method of managing an EX2500 switch does not provide a secure connection.

SSH is a protocol that enables remote administrators to log securely into the EX2500 over a network to execute management commands.

SSH provides the following benefits:

- Authentication of remote administrators
- Identifying the administrator using Name and Password
- Authorization of remote administrators
- Determining the permitted actions and customizing service for individual administrators
- Encryption of management messages
- Encrypting messages between the remote administrator and switch

The EX2500 implementation of SSH supports versions 1.0 and 2.0 and SSH client versions 1.5 through 2.x.

Configuring SSH Features on the Switch

SSH is disabled by default. Before you can use SSH commands on the switch, turn on SSH as follows:

```
ex2500(config)# ssh enable
```

SSH Encryption of Management Messages

The following encryption and authentication methods are supported for SSH:

- Server Host Authentication: Client RSA authenticates the switch at the beginning of every connection.
- Key Exchange: RSA.
- Encryption: 3DES-CBC and DES.
- User Authentication: Local password authentication.

Generating RSA Host and Server Keys for SSH Access

To support the SSH server feature, two sets of RSA keys (host and server keys) are required. The host key is 1024 bits and is used to identify the EX2500 switch. The server key is 768 bits and is used to make it impossible for someone to decipher a captured session by breaking into the EX2500 switch at a later time.

When the SSH server is first enabled and applied, the switch automatically generates the RSA host and server keys, which are stored in the Flash memory. To configure RSA host and server keys, enter the following commands to generate them manually:

```
ex2500(config)# ssh generate-host-key
ex2500(config)# ssh generate-server-key
```

When the switch reboots, it will retrieve the host and server keys from the Flash memory. If these two keys are not available in the flash and if the SSH server feature is enabled, the switch automatically generates them during the system reboot. This process might take several minutes to complete.

The switch can automatically regenerate the RSA server key. To set the interval of RSA server key autogeneration, use the following command:

```
ex2500(config)# ssh interval <number of hours (0-24)>
```

A value of 0 (zero) denotes that RSA server key autogeneration is disabled. When the interval value is greater than 0, the switch will autogenerate the RSA server key every specified interval. However, RSA server key generation is skipped if the switch is busy doing other key or cipher generation when the timer expires.



NOTE: The switch can perform only one session of key or cipher generation at a time. Thus, an SSH client will not be able to log in if the switch is performing key generation at that time, or if another client has logged in immediately prior. Also, key generation will fail if an SSH client is logging in at that time.

SSH Integration with RADIUS and TACACS+ Authentication

SSH is integrated with RADIUS authentication. After the RADIUS server is enabled on the switch, all subsequent SSH authentication requests will be redirected to the specified RADIUS servers for authentication. The redirection is transparent to the SSH clients.

SSH is integrated with TACACS+ authentication. After the TACACS+ server is enabled on the switch, all subsequent SSH authentication requests will be redirected to the specified TACACS+ servers for authentication. The redirection is transparent to the SSH clients.

End User Access Control

The EX2500 switch allows an administrator to define end user accounts that permit end users to perform operation tasks via the switch CLI commands. Once end user accounts are configured and enabled, the switch requires username-password authentication.

For example, an administrator can assign a user, who can then log in to the switch and perform operational commands (effective only until the next switch reboot).

Considerations for Configuring End User Accounts

- A maximum of 10 user IDs are supported on the switch.
- The EX2500 switch supports end user support for console, Telnet, EX2500 Web Device Manager, and SSHv1 or SSHv2 access to the switch.
- If RADIUS authentication is used, the user password on the RADIUS server will override the user password on the EX2500 switch. Also note that the password change command on the switch only modifies the use switch password and has no effect on the user password on the RADIUS server. RADIUS authentication and a user password cannot be used concurrently to access the switch.
- Passwords for end users can be up to 128 characters in length.

User Access Control

The end user access control commands allow you to configure end user accounts.

Setting Up User IDs

Up to 10 user IDs can be configured. Use the following commands to define usernames and passwords:

```
ex2500(config)# access user 1 name <1-8 characters>
ex2500(config)# access user 1 password
```

```
Changing user1 password; validation required:
Enter current admin password: <current administrator password>
Enter new user1 password: <new user password>
Re-enter new user1 password: <new user password>
New user1 password accepted.
```

Defining a User's Access Level

The end user is by default assigned to the user access level (also known as class of service, or COS). COS for all user accounts has global access to all resources except for User COS, which has access to view-only resources that the user owns. For more information, see Table 4 on page 13.

To change the user's level, select one of the following options:

```
ex2500(config)# access user 1 level { user | operator | administrator }
```

Enabling or Disabling a User

An end user account must be enabled before the switch recognizes and permits login under the account. Once enabled, the switch requires any user to enter both username and password.

```
ex2500(config)# access user 1 enable
ex2500(config)# no access user 1 enable
```

Listing Current Users

The following command displays defined user accounts and whether or not each user is currently logged in to the switch.

```
ex2500# show access user
```

Username:

```
user      - Enabled - offline
oper      - Disabled - offline
admin     - Always Enabled - online 1 session
```

Current User ID table:

```
1: name jane      , ena, cos user      , password valid, online 1 session
2: name john      , ena, cos user      , password valid, online 2 sessions
```

Logging In to an End User Account

Once an end user account is configured and enabled, the user can log in to the switch using the username-password combination. The level of switch access is determined by the COS established for the end user account.

Chapter 2

VLANs

This chapter describes network design and topology considerations for using Virtual Local Area Networks (VLANs). VLANs commonly are used to split up groups of network users into manageable broadcast domains, to create logical segmentation of workgroups, and to enforce security policies among logical segments. The following topics are discussed in this chapter:

- VLAN Overview on page 21
- VLANs and Port VLAN ID Numbers on page 22
- VLAN Tagging on page 23
- VLAN Topologies and Design Considerations on page 26
- Private VLANs on page 29



NOTE: VLANs can be configured from the CLI. See the information on VLAN configuration and port configuration in the *EX2500 Ethernet Switch Command Reference*.

VLAN Overview

Setting up virtual LANs (VLANs) is a way to segment networks to increase network flexibility without changing the physical network topology. With network segmentation, each switch port connects to a segment that is a single broadcast domain. When a switch port is configured to be a member of a VLAN, it is added to a group of ports (workgroup) that belong to one broadcast domain.

You group ports into broadcast domains by assigning them to the same VLAN. Frames received in one VLAN can be forwarded only within that VLAN, and multicast, broadcast, and unknown unicast frames are flooded only to ports in the same VLAN. The EX2500 switch supports jumbo frames, up to 9,216 bytes.

VLANs and Port VLAN ID Numbers

VLAN Numbers

The EX2500 switch supports up to 1024 VLANs per switch. Even though the maximum number of VLANs supported at any given time is 1024, each can be identified with any number between 1 and 4094. VLAN 1 is the default VLAN for the data ports. VLAN 4095 is used by the management network, which includes the management port.

Use the following CLI command to view VLAN information:

```
ex2500(config)# show vlan
```

VLAN	Name	Status	Ports
1	VLAN 1	ena	1-24, po1-po36
2	VLAN 2	dis	empty
4095	Mgmt VLAN	ena	MGMT

PVID Numbers

Each port in the switch has a configurable default VLAN number, known as its *PVID*. By default, the PVID for all non-management ports is set to 1, which correlates to the default VLAN ID. The PVID for each port can be configured to any VLAN number between 1 and 4094.

Use the following CLI commands to view and configure PVIDs:

- Port information:

```
ex2500# show interface information
```

Alias	Port	Tag	Edge	Ln	Fld	PVID	NAME	VLAN(s)
1	1	y	n	e	e	1	1	1
2	2	y	n	e	e	1	2	1
3	3	y	n	e	e	1	3	1
4	4	y	n	e	e	1	4	1
5	5	y	n	e	e	1	5	1
6	6	y	n	e	e	1	6	1
...
24	24	n	n	e	e	1	24	1
MGMT	MGMT	n	n	d	d	4095	MGMT	4095

= PVID is tagged.

- Port configuration:

```
ex2500(config)# interface port 7
ex2500(config-if)# pvid 7
```

Each port on the switch can belong to one or more VLANs, and each VLAN can have any number of switch ports in its membership. Any port that belongs to multiple VLANs, however, must have VLAN *tagging* enabled (see “VLAN Tagging” on page 23).

VLAN Tagging

EX2500 software supports 802.1Q VLAN *tagging*, providing standards-based VLAN support for Ethernet systems.

Tagging places the VLAN identifier in the frame header of a packet, allowing each port to belong to multiple VLANs. When you add a port to multiple VLANs, you also must enable tagging on that port.

Since tagging fundamentally changes the format of frames transmitted on a tagged port, you must carefully plan network designs to prevent tagged frames from being transmitted to devices that do not support 802.1Q VLAN tags, or devices where tagging is not enabled.

The following important terms are used with the 802.1Q tagging feature:

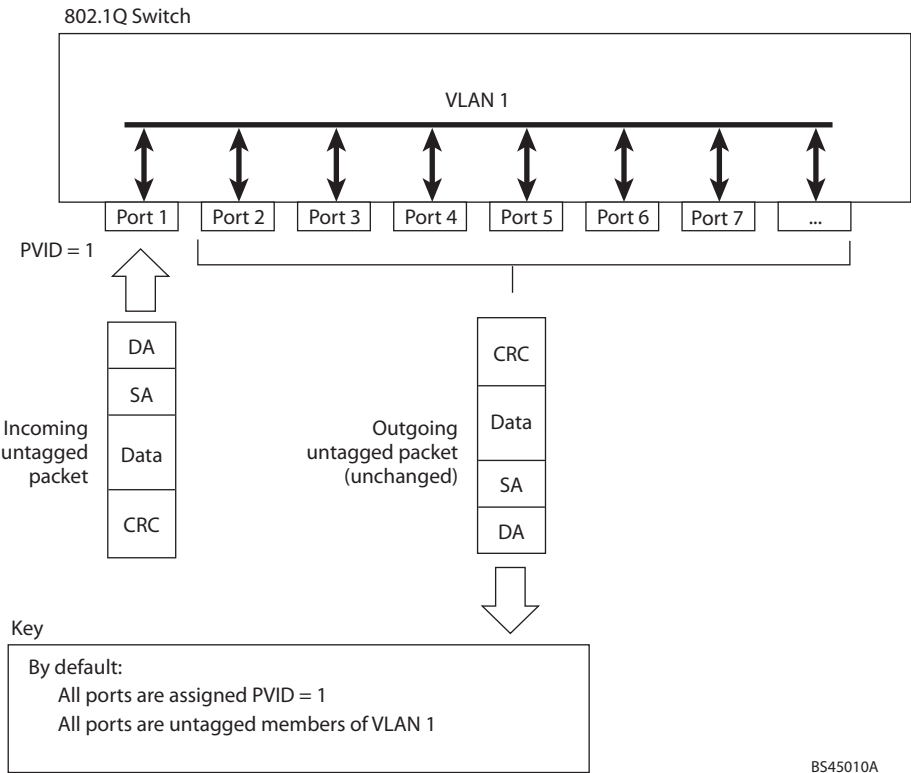
- VLAN identifier (VID)—The 12-bit portion of the VLAN tag in the frame header that identifies an explicit VLAN.
- Port VLAN identifier (PVID)—A classification mechanism that associates a port with a specific VLAN. For example, a port with a PVID of 3 (PVID = 3) assigns all untagged frames received on this port to VLAN 3. Any untagged frames received by the switch are classified with the PVID of the receiving port.
- Tagged frame—A frame that carries VLAN tagging information in the header. This VLAN tagging information is a 32-bit field (VLAN tag) in the frame header that identifies the frame as belonging to a specific VLAN. Untagged frames are marked (tagged) with this classification as they leave the switch through a port that is configured as a tagged port.
- Untagged frame—A frame that does not carry any VLAN tagging information in the frame header.
- Untagged member—A port that has been configured as an untagged member of a specific VLAN. When an untagged frame exits the switch through an untagged member port, the frame header remains unchanged. When a tagged frame exits the switch through an untagged member port, the tag is stripped and the tagged frame is changed to an untagged frame.
- Tagged member—A port that has been configured as a tagged member of a specific VLAN. When an untagged frame exits the switch through a tagged member port, the frame header is modified to include the 32-bit tag associated with the PVID. When a tagged frame exits the switch through a tagged member port, the frame header remains unchanged (original VID remains).



NOTE: If an 802.1Q tagged frame is received by a port that has VLAN tagging disabled and the port VLAN ID (PVID) is different from the VLAN ID of the packet, then the frame is dropped at the ingress port.

Figure 1 illustrates the default VLAN settings on the switch.

Figure 1: Default VLAN Settings



NOTE: The port numbers specified in these illustrations might not directly correspond to the physical port configuration of your switch model.

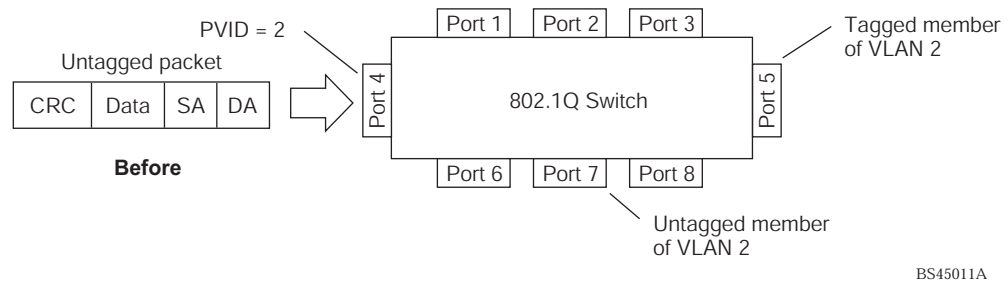
When a VLAN is configured, ports are added as members of the VLAN, and the ports are defined as either *tagged* or *untagged* (see Figure 2 through Figure 5).

The default configuration settings for the EX2500 switch have all ports set as untagged members of VLAN 1 with all ports configured as PVID = 1. In the default configuration example shown in Figure 1, all incoming packets are assigned to VLAN 1 by the default port VLAN identifier (PVID = 1).

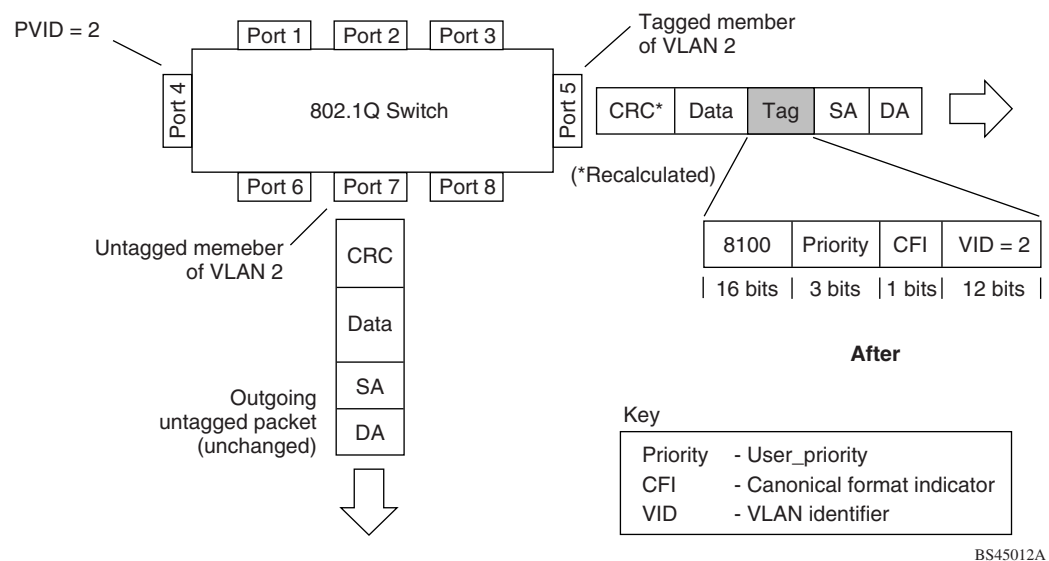
Figure 2 through Figure 5 illustrate generic examples of VLAN tagging. In Figure 2, untagged incoming packets are assigned directly to VLAN 2 (PVID = 2). Port 5 is configured as a *tagged* member of VLAN 2, and port 7 is configured as an *untagged* member of VLAN 2.




NOTE: The port assignments in the following figures are not meant to match the EX2500 switch.

Figure 2: Port-Based VLAN Assignment

As shown in Figure 3, the untagged packet is marked (tagged) as it leaves the switch through port 5, which is configured as a tagged member of VLAN 2. The untagged packet remains unchanged as it leaves the switch through port 7, which is configured as an untagged member of VLAN 2.

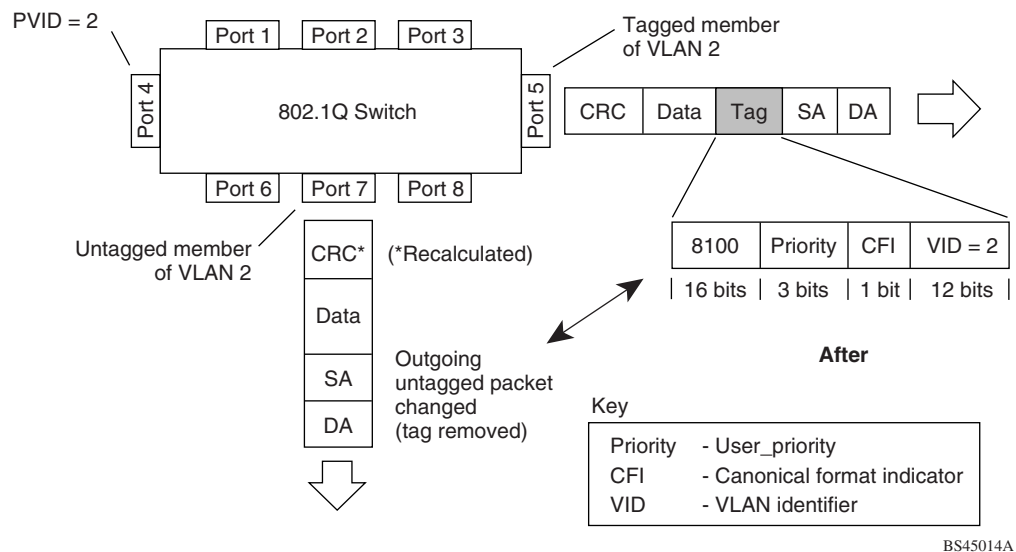
Figure 3: 802.1Q Tagging (after Port-Based VLAN Assignment)

In Figure 4, tagged incoming packets are assigned directly to VLAN 2 because of the tag assignment in the packet. Port 5 is configured as a *tagged* member of VLAN 2, and port 7 is configured as an *untagged* member of VLAN 2.

Figure 4: 802.1Q Tag Assignment

As shown in Figure 5, the tagged packet remains unchanged as it leaves the switch through port 5, which is configured as a tagged member of VLAN 2. However, the tagged packet is stripped (untagged) as it leaves the switch through port 7, which is configured as an untagged member of VLAN 2.

Figure 5: 802.1Q Tagging (after 802.1Q Tag Assignment)



VLAN Topologies and Design Considerations

This section discusses how you can connect users and segments to a host that supports many logical segments or subnets by using the flexibility of the multiple VLAN system. Be aware of the following VLAN properties on the EX2500 switch:

- By default, the EX2500 software is configured so that tagging is disabled on all ports.
- By default, the EX2500 software is configured so that all data ports are members of VLAN 1.
- By default, the EX2500 software is configured so that the management port is a member of VLAN 4095 (the management VLAN).
- If you configure Spanning Tree, note that Spanning Tree Groups 2-128 are allowed to contain only one VLAN.

VLAN Configuration Rules

VLANs operate according to specific configuration rules. When creating VLANs, consider the following rules that determine how the configured VLAN reacts in any network topology:


- All ports involved in trunking and port mirroring must have the same VLAN configuration. If a port is on a trunk with a mirroring port, the VLAN configuration cannot be changed. For more information trunk groups, see “Port Trunking Configuration Example” on page 45.

- All ports that are involved in port mirroring must have memberships in the same VLANs. If a port is configured for port mirroring, the port's VLAN membership cannot be changed. For more information on configuring port mirroring, see “Monitoring Ports with Port Mirroring” on page 81.

Multiple VLANs Configuration Example

Figure 6 shows a sample network consisting of an EX2500 switch configured with multiple VLANs with VLAN-tagged gigabit adapters.

Figure 6: Sample Network with Multiple VLANs



The multiple VLANs example in Figure 6 is described in Table 8.

Table 8: Components of Sample Network with Multiple VLANs (1 of 2)

Component	Description
EX2500 switch	This switch is configured with three VLANs that represent three different IP subnets. Five ports are connected downstream to servers. Two ports are connected upstream to routing switches. Uplink ports are members of all three VLANs, with VLAN tagging enabled.
Server 1	This server is a member of VLAN 1 and has presence in only one IP subnet. The associated switch port is only a member of VLAN 1, so tagging is disabled.
Server 2	This server is a member of VLAN 1 and has presence in only one IP subnet. The associated switch port is only a member of VLAN 1, so tagging is disabled.
Server 3	This server belongs to VLAN 2, and it is logically in the same IP subnet as Server 5. The associated switch port has tagging disabled.
Server 4	A member of VLAN 3, this server can communicate only with other servers via a router. The associated switch port has tagging disabled.

Table 8: Components of Sample Network with Multiple VLANs (2 of 2)

Component	Description
Server 5	A member of VLAN 1 and VLAN 2, this server can communicate only with Server 1, Server 2, and Server 3. The associated switch port has tagging enabled.
Enterprise Routing switches	These switches must have all three VLANs (VLAN 1, 2, 3) configured. They can communicate with Server 1, Server 2, and Server 5 via VLAN 1. They can communicate with Server 3 and Server 5 via VLAN 2. They can communicate with Server 4 via VLAN 3. Tagging on switch ports is enabled.



NOTE: VLAN tagging is required only on ports that are connected to other switches or on ports that connect to tag-capable end-stations, such as servers with VLAN-tagging adapters.

Use the following procedure to configure the sample network shown in Figure 6.

1. Enable VLAN tagging on server ports that support multiple VLANs.

```
ex2500(config)# interface port 5
ex2500(config-if)# tagging
ex2500(config-if)# exit
```

2. Enable tagging on uplink ports that support multiple VLANs.

```
ex2500(config)# interface port 19
ex2500(config-if)# tagging
ex2500(config-if)# exit
ex2500(config)# interface port 20
ex2500(config-if)# tagging
ex2500(config-if)# exit
```

3. Configure the VLANs and their member ports.

By default, all ports are members of VLAN 1, so configure only those ports that belong to other VLANs.

```
ex2500(config)# vlan 2
ex2500(config-vlan)# enable
ex2500(config-vlan)# member 3
ex2500(config-vlan)# member 5
ex2500(config-vlan)# member 19
ex2500(config-vlan)# member 20
ex2500(config-vlan)# exit
ex2500(config)# vlan 3
ex2500(config-vlan)# enable
ex2500(config-vlan)# member 4,19,20
ex2500(config-vlan)# exit
```

Private VLANs

Private VLANs provide Layer 2 isolation between the ports within the same broadcast domain. Private VLANs can control traffic within a VLAN domain, and provide port-based security for host servers.

Use private VLANs to partition a VLAN domain into sub-domains. Each sub-domain is comprised of one primary VLAN and one or more secondary VLANs, as follows:

- **Primary VLAN**—Carries unidirectional traffic downstream from promiscuous ports. Each private VLAN configuration has only one primary VLAN. All ports in the private VLAN are members of the primary VLAN.
- **Secondary VLAN**—Secondary VLANs are internal to a private VLAN domain, and are defined as follows:
 - **Isolated VLAN**—Carries unidirectional traffic upstream from the host servers toward ports in the primary VLAN and the gateway. Each private VLAN configuration can contain only one isolated VLAN.
 - **Community VLAN**—Carries upstream traffic from ports in the community VLAN to other ports in the same community, and to ports in the primary VLAN and the gateway. Each private VLAN configuration can contain multiple community VLANs.

After you define the primary VLAN and one or more secondary VLANs, you map the secondary VLAN(s) to the primary VLAN.

Private VLAN Ports

Private VLAN ports are defined as follows:

- **Promiscuous**—A promiscuous port is a port that belongs to the primary VLAN. The promiscuous port can communicate with all the interfaces, including ports in the secondary VLANs (isolated VLAN and community VLANs). Each promiscuous port can belong to only one private VLAN.
- **Isolated**—An isolated port is a host port that belongs to an isolated VLAN. Each isolated port has complete Layer 2 separation from other ports within the same private VLAN (including other isolated ports), except for the promiscuous ports.
 - Traffic sent to an isolated port is blocked by the private VLAN, except the traffic from promiscuous ports.
 - Traffic received from an isolated port is forwarded only to promiscuous ports.
- **Community**—A community port is a host port that belongs to a community VLAN. Community ports can communicate with other ports in the same community VLAN, and with promiscuous ports. These interfaces are isolated at Layer 2 from all other interfaces in other communities and from isolated ports within the private VLAN.

Private VLAN Configuration Guidelines

The following guidelines apply when configuring private VLANs:

- The default VLAN 1 cannot be a private VLAN.
- The management VLAN 4095 cannot be a private VLAN. The management port cannot be a member of a private VLAN.
- IGMP Snooping must be disabled on isolated VLANs.
- Each secondary port's (isolated port and community ports) PVID must match its corresponding secondary VLAN ID.
- Ports within a secondary VLAN cannot be members of other VLANs.
- All VLANs that make up the private VLAN must belong to the same Spanning Tree Group.

Private VLAN Configuration Example

Follow this procedure to configure a private VLAN.

1. Select a VLAN and define the private VLAN type as primary.

```
ex2500(config)# vlan 100
ex2500(config-vlan)# enable
ex2500(config-vlan)# member 2
ex2500(config-vlan)# private-vlan type primary
ex2500(config-vlan)# private-vlan enable
ex2500(config-vlan)# exit
```

2. Configure a secondary VLAN and map it to the primary VLAN.

```
ex2500(config)# vlan 110
ex2500(config-vlan)# enable
ex2500(config-vlan)# member 3
ex2500(config-vlan)# member 4
ex2500(config-vlan)# private-vlan type isolated
ex2500(config-vlan)# private-vlan map 100
ex2500(config-vlan)# private-vlan enable
ex2500(config-vlan)# exit
```

3. Verify the configuration.

```
ex2500(config)# show private-vlan
```

Private-VLAN	Type	Mapped-To	Status	Ports
100	primary	110	ena	2
110	isolated	100	ena	3-4

Chapter 3

Spanning Tree Protocol

When multiple paths exist on a network, Spanning Tree Protocol configures the network so that a switch uses only the most efficient path.

The following topics are discussed in this chapter:

- Spanning Tree Overview on page 31
- Rapid Spanning Tree Protocol on page 35
- Per VLAN Rapid Spanning Tree on page 37
- Multiple Spanning Tree Protocol on page 39
- Fast Uplink Convergence on page 41

Spanning Tree Overview

Spanning Tree Protocol (STP) detects and eliminates logical loops in a bridged or switched network. When multiple paths exist, Spanning Tree configures the network so that a switch uses only the most efficient path. If that path fails, Spanning Tree automatically sets up another active path on the network to sustain network operations.

The EX2500 switch supports the following Spanning Tree Protocols:

- IEEE 802.1D (2004) Rapid Spanning Tree Protocol (RSTP). RSTP enhances the Spanning Tree Protocol to provide rapid convergence on Spanning Tree Group 1.
- IEEE 802.1Q (2003) Multiple Spanning Tree Protocol (MSTP), which extends RSTP to multiple Spanning Tree Groups (STGs). MSTP provides both rapid convergence and load balancing in a VLAN environment, using multiple VLANs in each Spanning Tree Group (STG).
- Per VLAN Rapid Spanning Tree Plus Protocol (PVRST+), which enhances the RSTP protocol by adding the ability to have multiple Spanning Tree Groups. PVRST+ is based on IEEE 802.1w Rapid Spanning Tree Protocol.

The relationship between port, trunk groups, VLANs, and spanning trees is shown in Table 9.

Table 9: Ports, Trunk Groups, and VLANs

Switch Element	Belongs to
Port	Trunk group or One or more VLANs
Trunk group	One or more VLANs
VLAN (non-default)	RSTP: All VLANs in STG 1 PVRST + : One VLAN per Spanning Tree Group MSTP: Multiple VLANs per Spanning Tree Group



NOTE: Due to Spanning Tree's sequence of discarding, learning, and forwarding, lengthy delays might occur. You can use a port's **spanning-tree edge** command to permit a port that participates in Spanning Tree to bypass the Discarding and Learning states, and enter directly into the Forwarding state.

Bridge Protocol Data Units (BPDUs)

To create a spanning tree, the switch generates a configuration Bridge Protocol Data Unit (BPDU), which it then forwards out of its ports. All switches in the Layer 2 network participating in the spanning tree gather information about other switches in the network through an exchange of BPDUs.

A BPDU is a 64-byte packet that is sent out at a configurable interval, which is typically set for 2 seconds. The BPDU is used to establish a path, much like a "hello" packet in IP routing. BPDUs contain information about the transmitting bridge and its ports, including bridge MAC address, bridge priority, port priority, and path cost.

The generic action of a switch upon receiving a BPDU is to compare the received BPDU to its own BPDU that it will transmit. If the received BPDU is better than its own BPDU, it will replace its BPDU with the received BPDU. Then, the switch uses this information to block any necessary ports.

Determining the Path for Forwarding BPDUs

When determining which port to use for forwarding and which port to block, the EX2500 switch uses information in the BPDU, including each bridge ID. A technique based on the "lowest root cost" is then computed to determine the most efficient path for forwarding.

Bridge Priority

The bridge priority parameter controls which bridge on the network is the STG root bridge. To make one switch become the root bridge, configure the bridge priority lower than all other switches and bridges on your network. The lower the value, the higher the bridge priority. Use the following command to configure the spanning-tree bridge priority:

```
ex2500(config)# spanning-tree stp 1 bridge priority <0-61440, in steps of 4096>
```

Port Priority

The port priority helps determine which bridge port becomes the root or designated port. The case for the root port is when 2 switches are connected using a minimum of two links with the same path-cost. The case for the designated port is in a network topology that has multiple bridge ports with the same path-cost connected to a single segment—the port with the lowest port priority becomes the designated port for the segment. Use the following command to configure the spanning-tree port priority (Interface Port mode):

```
ex2500(config-if)# spanning-tree stp 1 priority <0-240, in steps of 16>
```

Port Path Cost

The port path cost assigns lower values to high-bandwidth ports, such as 10 Gigabit Ethernet, to encourage their use. The objective is to use the fastest links so that the route with the lowest cost is chosen. A value of 0 indicates that the default cost will be computed for an auto-negotiated link speed.

Use the following command to modify the port path cost:

```
ex2500(config-if)# spanning-tree stp 1 path-cost <0-200000000>
```

Spanning Tree Group Configuration Guidelines

This section provides important information on configuring Spanning Tree Groups (STGs):

Changing the Spanning Tree Mode

When the spanning-tree mode is changed (for example, RSTP to MSTP):

- You must reconfigure spanning-tree parameters for each STG, including VLAN assignment.
- If an STG in RSTP mode is disabled and then re-enabled, the Spanning Tree topology does not converge rapidly.

Assigning a VLAN to a Spanning Tree Group

- If no VLANs exist beyond the default VLAN 1 see “Creating a VLAN” on page 34 for information on adding ports to VLANs.
- Assign the VLAN to the STG with the following command:


```
ex2500(config-if)# spanning-tree stp 1 vlan <1-4094>
```
- If the association between the Spanning Tree Group (STG) and a VLAN is broken, the spanning-tree parameters are cleared. Reconfigure all the parameters for the STG.

- Each STG must have a VLAN assigned to it before it becomes functional. You cannot configure other STG settings until the VLAN is assigned. If the STG VLAN is unassigned, other configuration settings are cleared. Assign a VLAN and reconfigure the STG settings.



NOTE: To ensure proper operation with switches that use Cisco Per VLAN Spanning Tree (PVST+), you must either create a separate STG for each VLAN, or manually add all associated VLANs into a single STG.

Creating a VLAN

When you create a VLAN, that VLAN automatically belongs to STG 1, the default STG. You can assign the VLAN to another STG.

- Move a newly created VLAN to an existing STG by following this order:
 1. Create the VLAN.
 2. Enable the VLAN.
 3. Add the VLAN to an existing STG.
- VLANs must be contained *within* a single STG; a VLAN cannot span multiple STGs. By confining VLANs within a single STG, you avoid problems with Spanning Tree blocking ports and causing a loss of connectivity within the VLAN. When a VLAN spans multiple switches, we recommend that the VLAN remain within the same Spanning Tree Group (have the same STG ID) across all the switches.
- If ports are tagged, all tagged ports can belong to multiple STGs.
- A port cannot be added directly to an STG. First add the port to a VLAN, then add the VLAN to the STG.

Rules for VLAN Tagged Ports

- Tagged ports can belong to more than one STG.
- Untagged ports can belong to only one STG.

Adding and Removing Ports from STGs

- When you add a port to a VLAN that belongs to an STG, the VLAN's member port is added to the STG. However, if the port you are adding is an untagged port and is already a member of an STG, that port will be removed from this STG and added to the new STG. An untagged port cannot belong to more than one STG.

For example, assume that VLAN 2 belongs to STG 2. You add an untagged port (port 5) that belongs to STG 2 to VLAN 2. The port becomes a member of STG 2, and the switch displays a message to inform you that the PVID changed from 1 to 2:

Port 5 is an UNTAGGED port and its PVID changed from 1 to 2.

- When you remove a port from a VLAN that belongs to an STG, that port is removed from the STG. However, if that port belongs to another VLAN in the same STG, the port remains in the STG.

As an example, assume that port 1 belongs to VLAN 2, and VLAN 2 belongs to STG 2. When you remove port 1 from VLAN 2, port 1 is also removed from STG 2. The port moves to the default VLAN 1.

However, if port 1 belongs to both VLAN 1 and VLAN 2 and both VLANs belong to STG 1, removing port 1 from VLAN 2 does not remove port 1 from STG 1 because VLAN 1 is still a member of STG 1.

- An STG cannot be deleted, only disabled. If you disable the STG while it still contains VLAN members, Spanning Tree will be off on all ports belonging to that VLAN.

The relationship between port, trunk groups, VLANs, and spanning trees is shown in Table 9.

Rapid Spanning Tree Protocol

Rapid Spanning Tree Protocol (RSTP) provides rapid convergence of the spanning tree and provides for fast re-configuration critical for networks carrying delay-sensitive traffic such as voice and video. RSTP significantly reduces the time to reconfigure the active topology of the network when changes occur to the physical topology or its configuration parameters. RSTP reduces the bridged-LAN topology to a single spanning tree.

RSTP parameters are configured in Spanning Tree Group 1. Spanning Tree Groups 2 through 128 do not apply to RSTP. There are new STP parameters to support RSTP, and some values to existing parameters are different.

RSTP is compatible with devices that run 802.1D (1998) Spanning Tree Protocol. If the switch detects 802.1D (1998) BPDUs, it responds with 802.1D (1998)-compatible data units. RSTP is not compatible with Per VLAN Spanning Tree (PVST+) protocol.

Port State Changes

The port state controls the forwarding and learning processes of Spanning Tree. In RSTP, the port state has been consolidated to the following: discarding, learning, and forwarding. Table 10 compares the port states between 802.1D (1998) Spanning Tree and 802.1D (2004) Rapid Spanning Trees.

Table 10: RSTP vs. STP Port States

Operational Status	STP Port State	RSTP Port State
Enabled	Blocking	Discarding
Enabled	Listening	Discarding
Enabled	Learning	Learning
Enabled	Forwarding	Forwarding
Disabled	Disabled	Discarding

Port Type and Link Type

Spanning tree configuration includes the following parameters to support RSTP and MSTP: edge port and link type.

Edge Port

A port that does not connect to a bridge is called an *edge port*. Edge ports can start forwarding as soon as the link is up.

Edge ports do not take part in Spanning Tree, and should not receive BPDUs. A port with edge enabled is intended to be connected directly to a host. If a port with edge enabled does receive a BPDU, it immediately begins working as a normal port, and participates in Spanning Tree.

Link Type

The link type determines how the port behaves in regard to Rapid Spanning Tree. The link type corresponds to the duplex mode of the port. A full-duplex link is point-to-point (p2p), while a half-duplex link should be configured as **shared**. If you select **auto** as the link type, the port dynamically configures the link type.

RSTP Configuration Guidelines

This section provides important information about configuring Rapid Spanning Tree Groups:

- When RSTP is turned on, STP parameters apply only to Spanning Tree Group 1.
- When RSTP is turned on, Spanning Tree Groups 2 through 128 are turned off.
- When RSTP is turned on, all VLANs are moved to Spanning Tree Group 1.

RSTP Configuration Example

This section provides steps to configure Rapid Spanning Tree on the EX2500 switch, using the command-line interface (CLI).

Rapid Spanning Tree Protocol is the default setting on the EX2500 switch.

Use the following procedure to configure Rapid Spanning Tree.

1. Configure port and VLAN membership on the switch.
2. Set the spanning-tree mode to Rapid Spanning Tree.

```
ex2500(config)# spanning-tree mode rstp
```

Per VLAN Rapid Spanning Tree

Per VLAN Rapid Spanning Tree Plus Protocol (PVRST+) enhances the RSTP protocol by adding the ability to have multiple Spanning Tree Groups (STGs). PVRST+ is based on IEEE 802.1w Rapid Spanning Tree Protocol.

In PVRST mode, the EX2500 switch supports a maximum of 128 Spanning Tree Groups (STGs). Multiple STGs provide multiple data paths, which can be used for load balancing and redundancy.

To enable load balancing between two EX2500 switches using multiple STGs, configure each path with a different VLAN and then assign each VLAN to a separate STG. Each STG is independent. Each STG sends its own Bridge Protocol Data Units (BPDUs), and each STG must be configured independently.

The STG, or bridge group, forms a loop-free topology that includes one or more virtual LANs (VLANs). The switch supports 128 STGs running simultaneously. The default STG 1 may contain multiple VLANs. STGs 2 through 128 each may contain only one VLAN.

Default Spanning Tree Configuration

In the default configuration, a single STG (STG 1) includes all non-management ports on the switch. This is called the default STG. Although ports can be added to or deleted from the default STG, the default STG cannot be deleted from the system.

All other STGs, except the default STG 1, are empty and you must assign a VLAN to the STG. However, you cannot assign ports directly to an STG. Instead, you add ports to a VLAN and add the VLAN to the STG. Each STG is enabled by default, and assigned an ID number from 2 to 128.


By default, the spanning tree on the management port is turned off.

Why Do We Need Multiple Spanning Trees?

The following examples describe why we need multiple spanning trees.


In Figure 7, VLAN 1 and VLAN 2 pass traffic between switch 1 and switch 2. If you have a single Spanning Tree Group, the switches see an apparent physical loop, and one VLAN might become blocked, affecting connectivity, even though no logical loop exists. VLAN 2 traffic is blocked unnecessarily.

Figure 7: Two VLANs on One Spanning Tree Group



In Figure 8, VLAN 1 and VLAN 2 belong to different Spanning Tree Groups. The two instances of Spanning Tree separate the topology without forming a loop. Both VLANs can forward packets between the switches without losing connectivity.

Figure 8: Two VLANs, Each on a Different Spanning Tree Group



PVRST Configuration Guidelines

This section provides important information about configuring Per VLAN Rapid Spanning Tree Groups:

- By default, STGs 2 through 128 are empty, and STG 1 contains all configured VLANs until individual VLANs are assigned to other STGs. The EX2500 switch allows only one VLAN per STG, except for STG 1.
- If the ports are tagged, each port sends out a special BPDU containing the tagged information.
- When a tagged port belongs to more than one STG, the egress BPDUs are tagged to distinguish the BPDUs of one STG from those of another STG.

Configuring PVRST

This configuration shows how to configure PVRST+ on the switch.

1. Set the spanning-tree mode to PVRST+.

```
ex2500(config)# spanning-tree mode pvrst
```

2. Configure port membership for VLAN 1 and VLAN 2. Define the STGs for each VLAN.

By default, port 1 is a member of VLAN 1, and VLAN 1 is assigned to STG 1. Add port 2 to VLAN 2, and assign VLAN 2 to STG 2.

```
ex2500(config)# vlan 2
ex2500(config-vlan)# enable
ex2500(config-vlan)# member 2
ex2500(config-vlan)# stg 2
ex2500(config-vlan)# exit
```

Multiple Spanning Tree Protocol

Multiple Spanning Tree Protocol (MSTP) extends Rapid Spanning Tree Protocol through multiple Spanning Tree Groups, using multiple VLANs in each STG. MSTP supports up to 32 Spanning Tree instances, that correspond to Spanning Tree Groups 1 through 32. For more information about Spanning Tree Protocol, see “Spanning Tree Protocol” on page 31.

In Multiple Spanning Tree Protocol (MSTP), several VLANs can be mapped to each Spanning Tree instance. Each Spanning Tree instance is independent of other instances. MSTP allows frames assigned to different VLANs to follow separate paths, each path based on an independent Spanning Tree instance. This approach provides multiple forwarding paths for data traffic, enabling load balancing, and reducing the number of Spanning Tree instances required to support a large number of VLANs.

MSTP Region

A group of interconnected bridges that share the same attributes is called an MST region. Each bridge within the region must share the following attributes:

- Alphanumeric name
- Revision number
- VLAN-to-STG mapping scheme

MSTP provides rapid re-configuration, scalability, and control due to the support of regions, and support for multiple Spanning Tree instances within each region.

Common Internal Spanning Tree

The Common Internal Spanning Tree (CIST) provides a common form of Spanning Tree Protocol, with one Spanning Tree instance that can be used throughout the MSTP region. CIST allows the switch to interoperate with legacy equipment, including devices that run IEEE 802.1D (1998).

CIST allows the MSTP region to act as a virtual bridge to other bridges outside of the region, and provides a single Spanning Tree instance to interact with them.

CIST port configuration includes Hello time, path-cost, and interface priority. These parameters do not affect Spanning Tree Groups 1 through 32. They apply only when the CIST is used.


MSTP Configuration Guidelines

This section provides important information about configuring Multiple Spanning Tree Groups:

- When MSTP is turned on, the switch automatically moves all VLANs to the CIST. When MSTP is turned off, the switch moves all VLANs from the CIST to STG 1.
- When enabling MSTP, you must configure a Region Name, and a default version number of 0 (zero) is configured automatically. Each bridge in the region must have the same name, version number, and VLAN mapping.

Figure 9 shows how multiple spanning trees can provide redundancy without wasting any uplink ports. In this example, the server ports are split between two separate VLANs. Both VLANs belong to two different Multiple Spanning Tree (MSTP) Groups. The spanning-tree *priority* values are configured so that each routing switch is the root for a different MSTP instance. All of the uplinks are active, with each uplink port backing up the other.

Figure 9: Implementing Multiple Spanning Tree Groups



Multiple Spanning Tree Groups Configuration Example

This configuration shows how to configure MSTP Groups on the switch, as shown in Figure 9.

1. Configure port membership and define the Spanning Tree Groups (STGs) for VLAN 1.

Enable tagging on uplink ports that share VLANs. Port 19 and port 20 connect to the Enterprise Routing switches.

```
ex2500(config)# interface port 19
ex2500(config-if)# tagging
ex2500(config-if)# exit
ex2500(config)# interface port 20
ex2500(config-if)# tagging
ex2500(config-if)# exit
```

Add server ports 1 and 2 to VLAN 1. Add uplink ports 19 and port 20 to VLAN 1.

```
ex2500(config)# vlan 1
ex2500(config-vlan)# enable
ex2500(config-vlan)# member 1
ex2500(config-vlan)# member 2
ex2500(config-vlan)# member 19
ex2500(config-vlan)# member 20
ex2500(config-vlan)# stg 1
ex2500(config-vlan)# exit
```

2. Configure Multiple Spanning Tree Protocol.

Configure the MSTP region name and version, and set the spanning tree mode to mst.

```
ex2500(config)# spanning-tree mstp name MyRegion
ex2500(config)# spanning-tree mode mst
ex2500(config)# spanning-tree mstp version 100
```

3. Configure port membership and define the Spanning Tree Groups (STGs) for VLAN 2.

Add server ports 3, 4, and 5 to VLAN 2. Add uplink ports 19 and 20 to VLAN 2. Assign VLAN 2 to Spanning Tree Group 2.

```
ex2500(config)# vlan 2
ex2500(config-vlan)# enable
ex2500(config-vlan)# member 3
ex2500(config-vlan)# member 4
ex2500(config-vlan)# member 5
ex2500(config-vlan)# member 19
ex2500(config-vlan)# member 20
ex2500(config-vlan)# stg 2
ex2500(config-vlan)# exit
```



NOTE: Each Spanning Tree Group (STG) is enabled by default.

Fast Uplink Convergence

Fast Uplink Convergence enables the EX2500 switch to recover quickly from the failure of the primary link or trunk group in a Layer 2 network using Spanning Tree Protocol. Normal recovery can take as long as 50 seconds, while the backup link transitions from Blocking to Listening to Learning and then Forwarding states. With Fast Uplink Convergence enabled, the EX2500 switch immediately places the secondary path into Forwarding state, and sends multicasts of addresses in the forwarding database (FDB) and ARP table over the secondary link so that upstream switches can learn the new path.



NOTE: In order for Fast Uplink Convergence to be functional, the switch must be running in PVRST+ mode and must be linked to switches running STP or PVST.

Configuration Guidelines

When you enable Fast Uplink Convergence, the EX2500 switch automatically makes the following configuration changes:

- Sets the bridge priority to 61440 so that it does not become the root switch.
- Increases the cost of all ports by 30000, across all VLANs and Spanning Tree Groups. This ensures that traffic never flows through the EX2500 switch to get to another switch unless there is no other path.

These changes are reversed if the feature is disabled.

Configuring Fast Uplink Convergence

Use the following CLI command to enable Fast Uplink Convergence on all ports:

```
ex2500(config)# spanning-tree uplinkfast
```


Chapter 4

Ports and Trunking

Trunk groups can provide super-bandwidth, multi-link connections between switches or other trunk-capable devices. A trunk group is a group of ports that act together, combining their bandwidth to create a single, larger virtual link. This chapter provides configuration background and examples for trunking multiple ports together.

- Trunking Overview on page 43
- Port Trunking Configuration Example on page 45
- Configurable Trunk Hash Algorithm on page 47
- Link Aggregation Control Protocol on page 47



NOTE: Port trunking is also known as link aggregation.

Trunking Overview

When using port trunk groups between two switches, as shown in Figure 10 on page 45, you can create a virtual link between the switches, operating at up to 120 gigabits per second, depending on how many physical ports are combined. Each EX2500 switch supports up to 12 static trunk groups (portchannels) and up to 24 Link Aggregation Control Protocol (LACP trunk groups, consisting of 1 to 12 ports in each group.

Trunk groups are also useful for connecting a EX2500 switch to third-party devices that support link aggregation, such as routers and switches with EtherChannel technology (*not* ISL trunking technology) and Sun's Quad Fast Ethernet Adapter. Trunk group technology is compatible with these devices when they are configured manually.

Statistical Load Distribution

Network traffic is distributed statistically between the ports in a trunk group. The switch can use a combination of Layer 2 MAC and Layer 3 IP address information, present in each transmitted frame, to determine load distribution.

Each packet's particular MAC or IP address information results in selecting one line in the trunk group for data transmission. The more data streams are feeding the trunk lines, the more evenly traffic is distributed.

Built-In Fault Tolerance

Since each trunk group is made up of multiple physical links, the trunk group is inherently fault tolerant. As long as one connection between the switches is available, the trunk remains active.

Statistical load balancing is maintained whenever a port in a trunk group is lost or returned to service.

Before Configuring Static Trunks

When you create and enable a static trunk, the trunk members (switch ports) take on certain settings necessary for correct operation of the trunking feature.

Before you configure your trunk, you must consider these settings, along with specific configuration rules, as follows:

1. Read the configuration rules provided in the section, “Trunk Group Configuration Rules” on page 44.
2. Determine which switch ports (up to 12) are to become *trunk members* (the specific ports making up the trunk).

Ensure that the chosen switch ports are set to **enabled**. Trunk member ports must have the same VLAN and Spanning Tree configuration.

3. Consider how the existing Spanning Tree will react to the new trunk configuration. See “Spanning Tree Protocol” on page 31 for Spanning Tree Group configuration guidelines.
4. Consider how existing VLANs will be affected by the addition of a trunk.

Trunk Group Configuration Rules

The trunking feature operates according to specific configuration rules. When creating trunks, consider the following rules that determine how a trunk group reacts in any network topology:


- All trunks must originate from one device, and lead to one destination device.
- Any physical switch port can belong to only one trunk group.
- Trunking from third-party devices must comply with EtherChannel technology.
- When ports become members of a trunk, configuration parameters (except ACL and QoS) are applied per trunk. When a trunk group is formed, these parameters are configured for the trunk ID, which overrides the port-level parameters.
- All trunk member ports must be assigned to the same VLAN configuration before the trunk can be enabled.

- You cannot change the VLAN membership for a trunk group's member port. You can change the VLAN membership of the trunk group.
- When an active port is configured in a trunk, the port becomes a *trunk member* when you enable the trunk. The Spanning Tree parameters for the port then change to reflect the new trunk settings.
- All trunk members must be in the same Spanning Tree Group (STG) and can belong to only one Spanning Tree Group (STG). However if all ports are *tagged*, then all trunk ports can belong to multiple STGs.
- When a trunk is enabled, the trunk Spanning Tree participation setting takes precedence over that of any trunk member.
- You cannot configure a trunk member as a monitor port in a port-mirroring configuration.
- Trunks cannot be monitored by a monitor port; however, trunk members can be monitored.
- A trunk member cannot be configured as a monitor port.
- All ports in static trunks must have the same link configuration (speed, duplex, flow control).

Port Trunking Configuration Example

In the example shown in Figure 10, three ports are trunked between two switches.

Figure 10: Port Trunk Group Configuration Example



Use the following procedure to configure port trunking as shown in the Figure 10 example. You must first connect to the each switch's command line interface (CLI) as the administrator.



NOTE: For details about accessing and using any of the menu commands described in this example, see the *EX2500 Ethernet Switch Command Reference*.

1. Follow these steps on the EX2500 switch:

- a. Define a trunk group.

```
ex2500(config)# portchannel 3 member 2,9, 16
ex2500(config)# portchannel 3 enable
```

- b. Verify the configuration.

```
ex2500(config)# show portchannel 3
```

Examine the resulting information. If any settings are incorrect, make appropriate changes.

2. Repeat the process on the other switch.

```
ex2500(config)# portchannel 1 member 1,11,18
ex2500(config)# portchannel 1 enable
```

```
ex2500(config)# show portchannel 1
```

3. Connect the switch ports that will be members in the trunk group.

Trunk group 3 (on the EX2500 switch) is now connected to trunk group 1 (on the other switch).



NOTE: In this example, two EX2500 switches are used. If a third-party device supporting link aggregation is used (such as routers and switches with EtherChannel technology or Sun's Quad Fast Ethernet Adapter), trunk groups on the third-party device should be configured manually. Connection problems could arise during automatic trunk group negotiation on the third-party device.

4. Examine the trunking information on each switch.

```
ex2500# show portchannel information
PortChannel 3: Enabled
port state:
  2: STG  1 forwarding
  9: STG  1 forwarding
 16: STG  1 forwarding
```

Information about each port in each configured trunk group is displayed. Make sure that trunk groups consist of the expected ports and that each port is in the expected state. The following restrictions apply:

- Any physical switch port can belong to only one trunk group.
- Up to 12 ports can belong to the same trunk group.
- All ports in static trunks must have the same link configuration (speed, duplex, flow control).
- Trunking from third-party devices must comply with EtherChannel technology.

Configurable Trunk Hash Algorithm

This feature allows you to configure parameters for the trunk hash algorithm, instead of using the default values.

Use the IP Trunk Hash commands to configure new default behavior for Layer 2 traffic and Layer 3 traffic. The trunk hash settings affect both static trunks and LACP trunks.

You can select a minimum of one or a maximum of two parameters to create one of the following configurations:

- Source MAC (SMAC):
`ex2500(config)# portchannel hash source-mac-address`
- Destination MAC (DMAC):
`ex2500(config)# portchannel hash destination-mac-address`
- Source MAC (SMAC) + Destination MAC (DMAC):
`ex2500(config)# portchannel hash source-destination-mac`
- Source IP (SIP):
`ex2500(config)# portchannel hash source-ip-address`
- Destination IP (DIP):
`ex2500(config)# portchannel hash destination-ip-address`
- Source IP (SIP) + Destination IP (DIP):
`ex2500(config)# portchannel hash source-destination-ip`

Link Aggregation Control Protocol

Link Aggregation Control Protocol (LACP) is an IEEE 802.3ad standard for grouping several physical ports into one logical port (known as a dynamic trunk group or Link Aggregation group) with any device that supports the standard. Please see IEEE 802.3ad-2002 for a full description of the standard.

The 802.3ad standard allows standard Ethernet links to form a single Layer 2 link using the Link Aggregation Control Protocol (LACP). Link aggregation is a method of grouping physical link segments of the same media type and speed in full duplex, and treating them as if they were part of a single, logical link segment. If a link in an LACP trunk group fails, traffic is reassigned dynamically to the remaining link(s) of the dynamic trunk group.



NOTE: LACP implementation in the EX2500 switch does not support the Churn machine, an option used to detect if the port is operable within a bounded time period between the actor and the partner. Only the Marker Responder is implemented, and there is no marker protocol generator.

A port's Link Aggregation Identifier (LAG ID) determines how the port can be aggregated. The Link Aggregation ID (LAG ID) is constructed mainly from the *system ID* and the port's *admin key*, as follows:

- **System ID**—An integer value based on the switch's MAC address and the system priority assigned in the CLI.
- **Admin key**—An integer value (13-65535) for the port that you can configure in the CLI. Each switch port that participates in the same LACP trunk group must have the same *admin key* value. The admin key is *local significant*, which means the partner switch does not need to use the same admin key value.

For example, consider two switches, an Actor (the EX2500 switch) and a Partner (another switch), as shown in Table 11.

Table 11: Actor vs. Partner LACP Configuration

Actor Switch	Partner Switch 1
Port 7 (admin key = 100)	Port 1 (admin key = 50)
Port 8 (admin key = 100)	Port 2 (admin key = 50)

In the configuration shown in Table 11, Actor switch port 7 and port 8 aggregate to form an LACP trunk group with Partner switch port 1 and port 2.

LACP automatically determines which member links can be aggregated and then aggregates them. It provides for the controlled addition and removal of physical links for the link aggregation.

Each port on the switch can have one of the following LACP modes:

- **off** (default)—The user can configure this port in to a regular static trunk group.
- **active**—The port is capable of forming an LACP trunk. This port sends LACPDU packets to partner system ports.
- **passive**—The port is capable of forming an LACP trunk. This port responds only to the LACPDU packets sent from an LACP *active* port.

Each active LACP port transmits LACP data units (LACPDUs), while each passive LACP port listens for LACPDUs. During LACP negotiation, the admin key is exchanged. The LACP trunk group is enabled as long as the information matches at both ends of the link. If the admin key value changes for a port at either end of the link, that port's association with the LACP trunk group is lost.

When the system is initialized, all ports by default are in LACP off mode and are assigned unique admin keys. To make a group of ports aggregatable, you assign them all the same admin key. You must set the port's LACP mode to **active** to activate LACP negotiation. You can set other port's LACP mode to **passive**, to reduce the amount of LACPDU traffic at the initial trunk-forming stage.

Use the following command to check whether the ports are trunked:

```
ex2500# show lacp information
```

LACP Configuration Guidelines

Consider the following guidelines when you configure LACP trunks:

- On static trunks, configuration parameters (such as settings for VLAN membership, ACLs, and QoS) are applied per port. When an LACP trunk group is formed, these parameters are configured for the trunk ID, which overrides the port-level parameters.
- The range of potential LACP trunk IDs is 13 through 36.
- When an LACP trunk forms, the trunk ID is determined by the lowest port number in the trunk. For example, if the lowest port number is 1, then the LACP trunk ID is 13.
- The LACP trunk ID can change if the link is lost on the lowest port in the group. When the trunk ID changes, trunk-level parameters are cleared. To avoid losing configuration parameters, configure LACP trunk-level parameters for all possible trunk IDs.

Configuring LACP

Use the following procedure to configure LACP for port 7 and port 8 to participate in link aggregation.

1. Define the admin key on port 7. Only ports with the same admin key can form an LACP trunk group.

```
ex2500(config)# interface port 7-8
ex2500(config-if)# lacp key 100
```

2. Set the LACP mode.

```
ex2500(config-if)# lacp mode active
ex2500(config-if)# exit
```

Optionally Reducing LACP Timeout

The LACP timeout period is the number of seconds that elapse before the switch invalidates LACP data from a remote partner. The default LACP timeout value is **long** (90 seconds).

We recommend that you use the default **long** timeout to reduce LAPDU processing. If the CPU utilization rate of your switch remains at 100 % for periods of 90 seconds or more, consider using static trunks instead of LACP. However, if CPU use is low, you can set the LACP timeout value on the switch to **short** (3 seconds), instead.

```
ex2500(config-if)# lACP timeout short  
ex2500(config-if)# exit
```


Chapter 5

Quality of Service

Quality of Service features allow you to allocate network resources to mission-critical applications at the expense of applications that are less sensitive to such factors as time delays or network congestion. You can configure your network to prioritize specific types of traffic, ensuring that each type receives the appropriate Quality of Service (QoS) level.

The following topics are discussed in this chapter:


- QoS Overview on page 51
- Using ACL Filters on page 52
- Using Storm Control Filters on page 59
- Using DSCP Values to Provide QoS on page 60
- Using 802.1p Priority to Provide QoS on page 63
- Queuing and Scheduling on page 64

QoS Overview

QoS helps you allocate guaranteed bandwidth to the critical applications, and limit bandwidth for less critical applications. Applications such as video and voice must have a certain amount of bandwidth to work correctly; using QoS, you can provide that bandwidth when necessary. Also, you can put a high priority on applications that are sensitive to timing out or that cannot tolerate delay, by assigning their traffic to a high-priority queue.

By assigning QoS levels to traffic flows on your network, you can ensure that network resources are allocated where they are needed most. QoS features allow you to prioritize network traffic, thereby providing better service for selected applications.

Figure 11 on page 52 shows the basic QoS model used by the switch.

Figure 11: QoS Model

The basic QoS model works as follows:

- Classify traffic:
 - Read the DSCP value.
 - Read the 802.1p priority value.
 - Match ACL filter parameters.
- Perform actions:
 - Permit packets.
 - Deny packets.
 - Map the 802.1p priority to a COS queue.
 - Map the DSCP to a COS queue.
 - Set the number of COS queues (1 through 8).
- Queue and schedule traffic:
 - Place packets in one of the COS queues.
 - Schedule transmission based on the COS queue.

Using ACL Filters

Access Control Lists (ACLs) are filters that allow you to classify data packets according to a particular content in the packet header, such as the source address, destination address, source port number, destination port number, and others. Packet classifiers identify flows for more processing. Each filter defines the conditions that must match for inclusion in the filter, and also the actions that are performed when a match is made.

ACLs are used to control whether packets are forwarded or blocked at the switch ports. ACLs can provide basic security for access to the network. For example, you can use an ACL to permit one host to access a part of the network, and deny another host access to the same area.

Each ACL contains rules that define the matching criteria for data packets. The ACL checks each packet against its rules, to determine if there is a match. If the packet matches the ACL's rules, the ACL performs its configured action: either permit or deny the packet.

The EX2500 switch supports the following ACL types:

- MAC Extended ACLs
- IP Standard ACLs
- IP Extended ACLs

MAC Extended ACLs

The switch supports up to 127 MAC Extended ACLs, numbered from 1 through 127. Use MAC Extended ACLs to filter traffic using the following criteria:

- Source or destination MAC address
- VLAN
- Ethernet protocol
- User priority criteria

To create a MAC Extended ACL:

```
ex2500(config)# access-list mac extended 1
ex2500(config-ext-macl)#
```

To delete a MAC Extended ACL:

```
ex2500(config)# no access-list mac extended 1
ex2500(config)#
```

IP Standard ACLs

The switch supports up to 128 IP ACLs (standard and extended), numbered from 128 through 254. Use IP Standard ACLs to filter traffic using source IP address or network mask and destination IP address or network mask.

To create an IP Standard ACL:

```
ex2500(config)# access-list ip 128 standard
ex2500(config-std-nacl)#
```

To delete an IP Standard ACL:

```
ex2500(config)# no access-list ip 128 standard
ex2500(config)#
```

IP Extended ACLs

The switch supports up to 128 IP ACLs (standard and extended), numbered from 128 through 254. Use IP Extended ACLs to filter traffic using the following criteria:

- Source IP address or network mask
- Destination IP address or network mask
- IP protocol number or name as shown in Table 12
- TCP/UDP application ports, as shown in Table 13 on page 55
- TCP flags
- ICMP message code and type
- Type of Service (ToS) value
- DSCP value

To create an IP Extended ACL:

```
ex2500(config)# access-list ip 128 extended
ex2500(config-ext-nacl)#
```

To delete an IP Extended ACL:

```
ex2500(config)# no access-list ip 128 extended
ex2500(config)#
```

Table 12: Well-Known Protocol Types

Number	Protocol Name
1	icmp
4	ip
6	tcp
17	udp
89	ospf
103	pim

Table 13: Well-Known Application Ports

Number	TCP/UDP Application	Number	TCP/UDP Application	Number	TCP/UDP Application
20	ftp-data	79	finger	179	bgp
21	ftp	80	http	194	irc
22	ssh	109	pop2	220	imap3
23	telnet	110	pop3	389	ldap
25	smtp	111	sunrpc	443	https
37	time	119	nntp	520	rip
42	name	123	ntp	554	rtsp
43	whois	143	imap	1645, 1812	RADIUS
53	domain	144	news	1813	RADIUS accounting
69	tftp	161	snmp	1985	hsrp
70	gopher	162	snmptrap		

Understanding ACL Priority

Each ACL has a unique priority value, based on its number. The lower the ACL number, the higher the priority, so ACL 1 has the highest priority.

The priority value is used to decide which ACL rule to apply when a packet matches one or more ACLs. When an incoming packet matches the highest priority ACL, the ACL's configured action takes place. The other assigned ACLs are considered in numeric order, from lowest to highest.

In the following example, the switch considers ACL 128 before ACL 130 because ACL 128 has a higher priority. The order in which the ACLs are assigned to a port does not affect their priority.

Port 1 access group

ACL IP Extended 128:

TCP

Port number = 80

Action = permit

ACL IP Extended 129:

TCP

Port number = 23

Action = deny

ACL IP Extended 130:

TCP

Port number = less than 100

Action = permit

Assigning ACLs to a Port

Once you configure an ACL, you must assign the ACL to a port. Each port can accept multiple ACLs. Note that higher-priority ACLs are considered first, and their action takes precedence over lower-priority ACLs.

When you assign an ACL to a port, the ACL acts only upon ingress traffic, not egress traffic.

To assign an ACL to a port:

```
ex2500(config)# interface port 1
ex2500(config-if)# ip access-group 128 in
ex2500(config-if)# exit
```

To delete an ACL from a port:

```
ex2500(config)# interface port 1
ex2500(config-if)# no ip access-group 128 in
ex2500(config-if)# exit
```

Viewing ACL Statistics

ACL statistics display how many packets hit (matched) each ACL. Use ACL statistics to check filter performance, and debug the ACL filters. You must enable statistics for each ACL that you want to monitor. Use the following command to enable statistics for the ACL:

```
ex2500(config)# access-list ip standard 128 statistics
```

Use the following command to view ACL statistics:

```
ex2500(config)# show access-list counters
```

ACL Configuration Examples

ACL Example 1—Blocking Traffic to a Host

Use this configuration to block traffic to a specific host. All traffic that ingresses port 1 is denied if it is destined for the host at IP address 100.10.1.1.

1. Configure an Access Control List.

```
ex2500(config)# access-list ip 150 standard
ex2500(config-std-nacl)# deny any host 100.10.1.1
ex2500(config-std-nacl)# exit
```

2. Assign the ACL to port 1.

```
ex2500(config)# interface port 1
ex2500(config-if)# ip access-group 150 in
ex2500(config-if)# exit
```

3. Verify the configuration.

```
ex2500# show access-lists 1

Standard IP Access List 1
-----
Source IP address      : 0.0.0.0
Source IP address mask : 0.0.0.0
Destination IP address : 100.10.1.1
Destination IP address mask : 255.255.255.255
In Port List           : 1
Filter Action           : Deny
Status                  : InActive
```

ACL Example 2—Blocking Traffic from a Source to a Destination

Use this configuration to block traffic from a network destined for a specific host address. All traffic that ingresses port 10 with source IP from the class 100.10.1.0/24 and destination IP 200.20.2.2 is denied.

1. Configure an Access Control List.

```
ex2500(config)# access-list ip 160 standard
ex2500(config-std-nacl)# deny 100.10.1.0 255.255.255.0 host 200.20.2.2
ex2500(config-std-nacl)# exit
```

2. Assign the ACL to port 10.

```
ex2500(config)# interface port 10
ex2500(config-if)# ip access-group 160 in
ex2500(config-if)# exit
```

ACL Example 3—Blocking HTTP Traffic

Use this configuration to block HTTP traffic on a port.

1. Configure an Access Control List.

```
ex2500(config)# access-list ip 170 extended
ex2500(config-ext-nacl)# deny tcp any any eq 80
ex2500(config-ext-nacl)# exit
```

2. Add the ACL to a port.

```
ex2500(config)# interface port 12
ex2500(config-if)# ip access-group 170 in
ex2500(config-if)# exit
```

ACL Example 4—Blocking All Except Certain Packets

Use this configuration to block all traffic except traffic of certain types. HTTP/HTTPS, DHCP, and ARP packets are permitted on the port. All other traffic is denied.

1. Configure one IP ACL for each type of traffic that you want to permit.

```
ex2500(config)# access-list ip 200 extended
ex2500(config-ext-nacl)# permit tcp any any eq 80
ex2500(config-ext-nacl)# exit
ex2500(config)# access-list ip 210 extended
ex2500(config-ext-nacl)# permit tcp any any eq 443
ex2500(config-ext-nacl)# exit
ex2500(config)# access-list ip 220 extended
ex2500(config-ext-nacl)# permit udp any any eq 67
ex2500(config-ext-nacl)# exit
ex2500(config)# access-list ip 230 extended
ex2500(config-ext-nacl)# permit udp any any eq 68
ex2500(config-ext-nacl)# exit
```

2. Configure IP ACLs to deny all other traffic.

The ACLs that allow traffic must have a higher priority than the ACLs that deny all traffic.

```
ex2500(config)# access-list ip 240 extended
ex2500(config-ext-nacl)# deny tcp any any
ex2500(config-ext-nacl)# exit
ex2500(config)# access-list ip 245 extended
ex2500(config-ext-nacl)# deny udp any any
ex2500(config-ext-nacl)# exit
```

3. Configure one MAC ACL for each type of traffic that you want to permit (ARP).

```
ex2500(config)# access-list mac extended 10
ex2500(config-ext-macl)# permit any any arp
ex2500(config-ext-macl)# exit
```

4. Assign the ACLs to a port.

```
ex2500(config)# interface port 7
ex2500(config-if)# ip access-group 200 in
ex2500(config-if)# ip access-group 210 in
ex2500(config-if)# ip access-group 220 in
ex2500(config-if)# ip access-group 230 in
ex2500(config-if)# ip access-group 240 in
ex2500(config-if)# ip access-group 245 in
ex2500(config-if)# mac access-group 10 in
```


Using Storm Control Filters

The EX2500 switch provides filters that can limit the number of the following packet types transmitted by switch ports:

- Broadcast packets
- Multicast packets
- Unknown unicast packets (destination lookup failure)

Broadcast Storms

Excessive transmission of broadcast or multicast traffic can result in a broadcast storm. A broadcast storm can overwhelm your network with constant broadcast or multicast traffic, and degrade network performance. Common symptoms of a broadcast storm are slow network response times and network operations timing out.

Unicast packets whose destination MAC address is not in the Forwarding Database are *unknown unicasts*. When an unknown unicast is encountered, the switch handles it like a broadcast packet and floods it to all other ports in the VLAN (broadcast domain). A high rate of unknown unicast traffic can have the same negative effects as a broadcast storm.

Configuring Storm Control

Configure broadcast filters on each port that requires broadcast storm control. Set a threshold that defines the total number of broadcast packets transmitted, in megabits per second. When the threshold is reached, no more packets of the specified type are transmitted.

To filter broadcast packets on a port, use the following commands:

```
ex2500(config)# interface port 1
ex2500(config-if)# broadcast-threshold <packet rate (100-10000)>
ex2500(config-if)# exit
```

To filter multicast packets on a port, use the following commands:

```
ex2500(config)# interface port 1
ex2500(config-if)# multicast-threshold <packet rate (100-10000)>
ex2500(config-if)# exit
```

To filter unknown unicast packets on a port, use the following commands:

```
ex2500(config)# interface port 1
ex2500(config-if)# dest-lookup-threshold <packet rate (100-10000)>
ex2500(config-if)# exit
```



NOTE: You can filter unknown unicast packets on no more than 16 ports at a time.

Using DSCP Values to Provide QoS

The switch uses the Differentiated Services (DiffServ) architecture to provide QoS functions. DiffServ is described in IETF RFCs 2474 and 2475.


The six most significant bits in the ToS byte of the IP header are defined as DiffServ Code Points (DSCP). Packets are marked with a certain value depending on the type of treatment the packet must receive in the network device. DSCP is a measure of the Quality of Service (QoS) level of the packet.

The switch can classify traffic by reading the DiffServ Code Point (DSCP) or IEEE 802.1p priority value, or by using filters to match specific criteria. When network traffic attributes match those specified in a traffic pattern, the policy instructs the switch to perform specified actions on each packet that passes through it. The packets are assigned to different Class of Service (COS) queues and scheduled for transmission.

Differentiated Services Concepts

To differentiate between traffic flows, packets can be classified by their DSCP value. As shown in Figure 12, the Differentiated Services (DS) field in the IP header is an octet, and the first six bits, called the DS Code Point (DSCP), can provide QoS functions. Each packet carries its own QoS state in the DSCP. There are 64 possible DSCP values (0 through 63).

Figure 12: Layer 3 IPv4 Packet



The switch can perform the following actions to the DSCP:

- Read the DSCP value of ingress packets.
- Map the DSCP value to a Class of Service queue (COSq).

The switch can use the DSCP value to direct traffic prioritization.

With DiffServ, you can establish policies to direct traffic. A policy is a traffic-controlling mechanism that monitors the characteristics of the traffic (for example, its source, destination, and protocol), and performs a controlling action on the traffic when certain characteristics are matched.

Per Hop Behavior

The DSCP value determines the Per Hop Behavior (PHB) of each packet. The PHB is the forwarding treatment given to packets at each hop. QoS policies are built by the application of a set of rules to packets, based on the DSCP value, as they hop through the network.

The default settings are based on the following standard PHBs, as defined in the IEEE standards:

- Expedited Forwarding (EF)—This PHB has the highest egress priority and lowest drop precedence level. EF traffic is forwarded ahead of all other traffic. EF PHB is described in RFC 2598.
- Assured Forwarding (AF)—This PHB contains four service levels, each with a different drop precedence, as shown in Table 14. Routers use drop precedence to determine which packets to discard last when the network becomes congested. AF PHB is described in RFC 2597.

Table 14: Assured Forwarding

Drop Precedence	Class 1	Class 2	Class 3	Class 4
Low	AF11 (DSCP 10)	AF21 (DSCP 18)	AF31 (DSCP 26)	AF41 (DSCP 34)
Medium	AF12 (DSCP 12)	AF22 (DSCP 20)	AF32 (DSCP 28)	AF42 (DSCP 36)
High	AF13 (DSCP 14)	AF23 (DSCP 22)	AF33 (DSCP 30)	AF43 (DSCP 38)

- Class Selector (CS)—This PHB has eight priority classes, with CS7 representing the highest priority, and CS0 representing the lowest priority, as shown in Table 15. CS PHB is described in RFC 2474.

Table 15: Class Selector

Priority	Class Selector	DSCP
Highest	CS7	56
	CS6	48
	CS5	40
	CS4	32
	CS3	24
	CS2	16
	CS1	8
Lowest	CS0	0

QoS Levels

Table 16 shows the default service levels provided by the switch, listed from highest to lowest importance.

Table 16: Default QoS Service Levels

Service Level	Default PHB	802.1p Priority
Critical	CS7	7
Network Control	CS6	6
Premium	EF, CS5	5
Platinum	AF41, AF42, AF43, CS4	4
Gold	AF31, AF32, AF33, CS3	3
Silver	AF21, AF22, AF23, CS2	2
Bronze	AF11, AF12, AF13, CS1	1
Standard	DF, CS0	0

DSCP Mapping

The switch can use the DSCP value of ingress packets to set the COS queue. Use the following command to view the default settings:

```
ex2500(config)# show qos dscp
  DSCP    CoS Queue
-----
    0      0
    1      0
    2      0
    3      0
    4      0
    5      0
    6      0
    7      0
    8      1
    9      1
   10      1
...
   54      6
   55      6
   56      7
   57      7
   58      7
   59      7
   60      7
   61      7
   62      7
   63      7
```

Use the following command to turn on DSCP re-marking globally:

```
ex2500# qos dscp enable
```

Use the following command to perform DSCP mapping:

```
ex2500# qos dscp transmit-queue <DSCP value (0-63)> <COSq (0-7)>
```


Using 802.1p Priority to Provide QoS

The EX2500 switch provides Quality of Service (QoS) functions based on the priority bits in a packet's VLAN header. (The priority bits are defined by the 802.1p standard within the IEEE 802.1Q VLAN header.) The 802.1p bits, if present in the packet, specify the priority that should be given to packets during forwarding. Packets with a numerically higher (non-zero) priority are given forwarding preference over packets with lower priority value.

The IEEE 802.1p standard uses eight levels of priority (0 through 7). Priority 7 is assigned to highest-priority network traffic, such as OSPF or RIP routing table updates, priorities 5 through 6 are assigned to delay-sensitive applications such as voice and video, and lower priorities are assigned to standard applications. A value of 0 (zero) indicates a "best effort" traffic prioritization, and this is the default when traffic priority has not been configured on your network. The switch can filter packets based on the 802.1p values.

Figure 13 shows the priority bits in a VLAN-tagged packet.

Figure 13: Layer 2 802.1q/802.1p VLAN-Tagged Packet



Ingress packets receive a priority value, as follows:

- **Tagged packets**—The switch reads the 802.1p priority in the VLAN tag.
- **Untagged packets**—The switch tags the packet and assigns an 802.1p priority value, based on the port's default 802.1p priority.

Egress packets are placed in a Class of Service (COS) queue based on the priority value, and scheduled for transmission based on the COS queue number. Higher COS queue numbers provide forwarding precedence.

The following is an example of 802.1p configuration:

1. Configure a port's default 802.1p priority value to 2.

```
ex2500(config)# interface port 1
ex2500(config-if)# dot1p 2
ex2500(config-if)# exit
```

2. Map the 802.1p priority value to a COS queue.

```
ex2500(config)# qos transmit-queue mapping 1 0
```

Queuing and Scheduling

The EX2500 switch has eight output Class of Service (COS) queues per port, into which each packet is placed. Each packet's 802.1p priority determines its COS queue. Higher COS queue numbers provide forwarding precedence.

You can map 802.1p priority value to a COS queue, as follows:

```
ex2500(config)# qos transmit-queue mapping <802.1p priority value (0-7)>  
<COS queue (0-7)>
```

Chapter 6

Remote Monitoring

Remote Monitoring (RMON) allows network devices to exchange network monitoring data. The following topics are discussed in this chapter:

- RMON Overview on page 65
- RMON Group 1—Statistics on page 66
- RMON Group 2—History on page 67
- RMON Group 3—Alarms on page 68
- RMON Group 9—Events on page 69

RMON Overview

RMON allows the switch to track events and trigger alarms when a threshold is reached and to notify administrators by issuing a syslog message or SNMP trap.

The RMON MIB provides an interface between the RMON agent on the switch and an RMON management application. The RMON MIB is described in RFC 1757.

The RMON standard defines objects that are suitable for the management of Ethernet networks. The RMON agent continuously collects statistics and proactively monitors switch performance. RMON allows you to monitor traffic flowing through the switch.

The switch supports the following RMON Groups, as described in RFC 1757:

- Group 1: Statistics
- Group 2: History
- Group 3: Alarms
- Group 9: Events

RMON Group 1—Statistics

The switch supports collection of Ethernet statistics as outlined in the RMON statistics MIB, in reference to **etherStatsTable**. You can configure RMON statistics on a per-port basis.

RMON statistics are sampled every second, and new data overwrites any old data on a given port.



NOTE: You must configure RMON statistics for the port before you can view RMON statistics.

Use the following procedure to configure RMON statistics:

1. Enable RMON on a port.

```
ex2500(config)# interface port 1
ex2500(config-if)# rmon enable
```

2. Configure the RMON statistics on a port.

```
ex2500(config)# interface port 1
ex2500(config-if)# rmon collection-stats 1
ex2500(config-if)# rmon collection-stats owner "port 1 rmon"
```

This configuration enables RMON statistics on port 1.

3. View RMON statistics for the port.

```
ex2500(config)# show rmon statistics
RMON is enabled
Collection 1 on 7 is active, and owned by port 1 rmon
Monitors ifEntry.1.7 which has
Received 0 octets, 0 packets,
0 broadcast and 0 multicast packets,
0 undersized and 0 oversized packets,
0 fragments and 0 jabbers,
0 CRC alignment errors and 0 collisions.
# of packets received of length (in octets):
64: 0, 65-127: 0, 128-255: 0,
256-511: 0, 512-1023: 0, 1024-1518: 0
```


RMON Group 2—History

The RMON History Group allows you to sample and archive Ethernet statistics for a specific interface during a specific time interval. History sampling is done per port.



NOTE: RMON port statistics must be enabled for the port before an RMON History Group can monitor the port.

Data is gathered during discreet sampling intervals and stored in “*buckets*.” At each configured interval, the History index takes a sample of the current Ethernet statistics and places them into a bucket. History data buckets reside in dynamic memory. When the switch is rebooted, the buckets are emptied.

Requested buckets are the number of buckets, or data slots, requested by the user for each History Group. *Granted buckets* are the number of buckets granted by the system, based on the amount of system memory available. The system grants a maximum of 50 buckets.

You can use an SNMP browser to view History samples, or use the following command:

```
ex2500# show rmon history
```

History MIB Object ID

The type of data that can be sampled must be of an `ifIndex` object type, as described in RFC 1213 and RFC 1573. The most common data type for the History sample is as follows:

```
1.3.6.1.2.1.2.2.1.1.x
```

The last digit (x) represents the number of the port to monitor. In the CLI, you do not need to enter the History MIB Object Identifier (OID), because the port is specified when you enter Interface Port mode.

Configuring RMON History

Perform the following steps to configure RMON History on a port.

1. Enable RMON on a port.

```
ex2500(config)# interface port 1
ex2500(config)# rmon enable
```

2. Configure the RMON History parameters for a port.

```
ex2500(config-if)# rmon collection-history 1 buckets 30
ex2500(config-if)# rmon collection-history 1 interval 120
ex2500(config-if)# rmon collection-history 1 owner "rmon port 1 history"
```

This configuration enables RMON History collection on port 1.

3. View RMON history for the port.

```
ex2500(config)# show rmon history
RMON is enabled
```

Index	IFOID	Interval	Rbnum	Gbnum
1	ifEntry.1.7	120	30	30

```
History Ether table is empty
```

RMON Group 3—Alarms

The RMON Alarm Group allows you to define a set of thresholds to determine network performance. When a configured threshold is crossed, an alarm is generated. For example, you can configure the switch to issue an alarm if more than 1,000 CRC errors occur during a 10-minute time interval.

Each Alarm index consists of a variable to monitor, a sampling time interval, and parameters for rising and falling thresholds. The Alarm Group can be used to track rising or falling values for a MIB object. The object must be a counter, gauge, integer, or time interval.

Use one of the following commands to correlate an Alarm index to an Event index:

```
ex2500(config)# rmon alarm <alarm number> rise-event <event number>
ex2500(config)# rmon alarm <alarm number> fall-event <event number>
```

When the alarm threshold is reached, the corresponding event is triggered.

Alarm MIB Objects

The most common data types used for alarm monitoring are **ifStats**: errors, drops, invalid CRCs, and so on. These MIB Object Identifiers (OIDs) correlate to the ones tracked by the History Group. An example of an ICMP statistic is as follows:

1.3.6.1.2.1.5.1.0 – **mgmt.icmp.icmplnMsgs**

This value represents the alarm's MIB OID, as a string. Note that for non-tables, you must supply a .0 to specify end node.

Configuring RMON Alarms

Configure the RMON Alarm parameters to track ICMP messages.

```
ex2500(config)# rmon alarm 1 oid 1.3.6.1.2.1.5.8.0 alarm-type rising rise-event 110
ex2500(config)# rmon alarm 1 interval-time 60
ex2500(config)# rmon alarm 1 rising-threshold 200
ex2500(config)# rmon alarm 1 sample-type delta
ex2500(config)# rmon alarm 1 owner "Alarm for icmplnEchos"
```

This configuration creates an RMON alarm that checks **icmplnEchos** on the switch once every minute. If the statistic exceeds 200 within a 60-second interval, an alarm is generated that triggers event index 110.

RMON Group 9—Events

The RMON Event Group allows you to define events that are triggered by alarms. An event can be a log message, an SNMP trap, or both.

When an alarm is generated, it triggers a corresponding event notification. Use the following commands to correlate an Event index to an alarm:

```
ex2500(config)# rmon alarm <alarm number> rise-event <event number>  
ex2500(config)# rmon alarm <alarm number> fall-event <event number>
```

RMON events use SNMP and syslogs to send notifications. Therefore, an SNMP trap host must be configured for trap event notification to work properly.

RMON uses a syslog host to send syslog messages. Therefore, an existing syslog host must be configured for event log notification to work properly. Each log event generates a syslog of type RMON that corresponds to the event.

Use the following commands to configure RMON event parameters:

```
ex2500(config)# rmon event 110 type log-only  
ex2500(config)# rmon event 110 description "SYSLOG_this_alarm"  
ex2500(config)# rmon event 110 owner "log icmplnEchos alarm"
```

This configuration creates an RMON event that sends a syslog message each time it is triggered by an alarm.

Chapter 7

IGMP

Internet Group Management Protocol (IGMP) is used by IP Multicast routers to learn about the existence of host group members on their directly attached subnet (see RFC 2236). The IP Multicast routers get this information by broadcasting IGMP Membership Queries and listening for IP hosts reporting their host group memberships. This process is used to set up a client/server relationship between an IP Multicast source that provides the data streams and the clients that want to receive the data.

The EX2500 switch can perform IGMP Snooping, and connect to static multicast routers (Mrouter).

The following topics are discussed in this chapter:

- IGMP Snooping on page 71
- FastLeave on page 72
- IGMPv3 Snooping on page 73
- IGMP Snooping Configuration Example on page 73
- Static Multicast Router on page 74

IGMP Snooping

IGMP snooping allows the switch to forward multicast traffic to only those ports that request it. IGMP Snooping prevents multicast traffic from being flooded to all ports. The switch learns which server hosts are interested in receiving multicast traffic, and forwards it only to ports connected to those servers.

IGMP snooping conserves bandwidth. With IGMP Snooping, the switch learns which ports are interested in receiving multicast data, and forwards multicast data to those ports only. In this way, other ports are not burdened with unwanted multicast traffic.

The switch can sense IGMP Membership Reports from attached clients and can act as a proxy to set up a dedicated path between the requesting host and a local IP Multicast router. After the pathway is established, the switch blocks the IP Multicast stream from flowing through any port that does not connect to a host member, thus conserving bandwidth.

The client-server path is set up as follows:

1. An IP Multicast router (Mrouter) sends *Membership Queries* to the switch, which forwards them to all ports in a given VLAN.
2. Hosts that want to receive the multicast data stream send *Membership Reports* to the switch, which sends a proxy Membership Report to the Mrouter.
3. The switch sets up a path between the Mrouter and the host, and blocks all other ports from receiving the multicast.
4. Periodically, the Mrouter sends Membership Queries to ensure that the host wants to continue receiving the multicast. If a host fails to respond with a Membership Report, the Mrouter stops sending the multicast to that path.
5. The host can send a *Leave Report* to the switch, which sends a proxy Leave Report to the Mrouter. The multicast path is terminated immediately.

The EX2500 switch supports the following IGMP capacities:

- IGMP versions 1, 2, and 3
- 1024 VLANs
- 128 M routers
- 1024 multicast groups



NOTE: Unknown multicast traffic is sent to all ports if the flood option is disabled. To enable or disable IGMP flood, use the following command:
ex2500(config)# **[no] ip igmp flood**

FastLeave

In normal IGMP operation, when the switch receives an IGMPv2 leave message, it sends a Group-Specific Query to determine if any other devices in the same group (and on the same port) are still interested in the specified multicast group traffic. The switch removes the affiliated port from that particular group, if the following conditions apply:

- If the switch does not receive an IGMP Membership Report within the query-response-interval.
- If no multicast routers have been learned on that port.

With FastLeave enabled on the VLAN, a port can be removed immediately from the port list of the group entry when the IGMP Leave message is received, unless a multicast router was learned on the port.

Enable FastLeave only on VLANs that have only one host connected to each physical port. To enable FastLeave, use the following command:

```
ex2500(config)# ip igmp fastleave <VLAN number (1-4094)>
```

IGMPv3 Snooping

IGMPv3 includes new membership report messages to extend IGMP functionality. The switch provides snooping capability for all types of IGMP version 3 (IGMPv3) Membership Reports.

IGMPv3 supports Source-Specific Multicast (SSM). SSM identifies session traffic by both source and group addresses.

The IGMPv3 implementation keeps records on the multicast hosts present in the network. If a host is already registered, when it sends an IS_INC, TO_INC, IS_EXC, or TO_EXC report, the switch overwrites the existing (port-host-group) registration with the new registration; the registrations of other hosts on the same group, same port are not changed. IS_INCLUDE and TO_INCLUDE reports with no source are not registered.

The switch supports the following IGMPv3 filter modes:

- **INCLUDE** mode—The host requests membership to a multicast group and provides a list of IP addresses from which it wants to receive traffic.
- **EXCLUDE** mode—The host requests membership to a multicast group and provides a list of IP addresses from which it *does not* want to receive traffic. This indicates that the host wants to receive traffic only from sources that are not part of the Exclude list. To disable snooping on EXCLUDE mode reports, use the following command:

```
ex2500(config)# no ip igmp snoop igmpv3 exclude
```

By default, the switch snoops the first eight sources listed in the IGMPv3 Group Record. Use the following command to change the number of snooping sources:

```
ex2500(config)# ip igmp snoop igmpv3 sources <1-64>
```

IGMPv3 Snooping is compatible with IGMPv1 and IGMPv2 Snooping. You can disable snooping on version 1 and version 2 reports, with the following command:

```
ex2500(config)# no ip igmp snoop igmpv3 v1v2
```

IGMP Snooping Configuration Example

This section provides steps to configure IGMP Snooping on the switch. Use the following procedure to configure IGMP Snooping:

1. Configure port and VLAN membership on the switch.
2. Enable IGMP Snooping.

```
ex2500(config)# ip igmp snoop enable
```

3. Add VLANs to IGMP Snooping.

```
ex2500(config)# ip igmp snoop vlan 1
```

4. Enable IGMPv3 Snooping (optional).

```
ex2500(config)# ip igmp snoop igmpv3 enable
```

5. View dynamic IGMP information.

```
ex2500# show ip igmp groups
```

Note: Local groups (224.0.0.x) are not snooped/relayed and will not appear.

Source	Group	VLAN	Port	Version	Mode	Expires	Fwd
10.1.1.1	232.1.1.1	2	4	V3	INC	4:16	Yes
10.1.1.5	232.1.1.1	2	4	V3	INC	4:16	Yes
*	232.1.1.1	2	4	V3	INC	-	No
10.10.10.43	235.0.0.1	9	1	V3	INC	2:26	Yes
*	236.0.0.1	9	1	V3	EXC	-	Yes

```
# show ip igmp mrouter
```

VLAN	Port	Version	Expires	Max Query Resp. Time	QRV	QQIC
1	4	V2	static	-	-	-
2	3	V3	4:09	128	2	125

These commands display information about IGMP Groups and Mrouters learned by the switch.

Static Multicast Router

A static multicast router (Mrouter) can be configured for a particular port on a particular VLAN. A static Mrouter does not have to be learned through IGMP Snooping. Any data port can accept a static Mrouter.

When you configure a static Mrouter on a VLAN, it replaces any dynamic Mrouters learned through IGMP Snooping.

Use the following procedure to configure a static multicast router (Mrouter):

1. For each Mrouter, configure a port or trunk group (1 through 24 or po1 through po24), VLAN (1 through 4094), and version (1 through 3).

```
ex2500(config)# ip igmp mrouter 5 1 2
```

The IGMP version is set for each VLAN, and cannot be configured separately for each Mrouter.

2. Verify the configuration.

```
ex2500(config)# show ip igmp mrouter
```


Chapter 8

High Availability Through Uplink Failure Detection

This chapter describes how to use Uplink Failure Detection (UFD) to ensure that network resources remain available if one switch is removed for service.

The following topics are discussed in this chapter:

- High Availability Overview on page 75
- Failure Detection Pair on page 76
- Spanning Tree Protocol with UFD on page 76
- UFD Configuration Guidelines on page 76
- UFD Configuration Example on page 77
- Monitoring UFD on page 77

High Availability Overview


The EX2500 switch supports high-availability network topologies.

Uplink Failure Detection (UFD) is designed to support Network Adapter Teaming. Network Adapter Teaming allows all the network interface cards (NICs) on each server to share the same IP address. The NICs are configured into a team. One NIC is the primary link, and the other is a standby link.

UFD allows the switch to monitor specific ports (Link to Monitor ports) to detect link failures. When the switch detects a link failure, it automatically disables specific ports (Link to Disable ports). Each corresponding server's network adapter can detect the disabled port and trigger a network-adapter failover to another port on the switch.

The switch automatically enables the control ports when the monitor ports return to service.

Figure 14 on page 76 shows a basic UFD configuration, with a Failure Detection Pair (FDP) that consists of one LtM (Link to Monitor) and one LtD (Link to Disable). When the switch detects a link failure in the LtM, it disables the ports in the LtD. The servers detect the disabled ports, which triggers a NIC failover.

Figure 14: Uplink Failure Detection Example

Failure Detection Pair

To use UFD, you must configure a Failure Detection Pair and then turn UFD on. A Failure Detection Pair consists of the following groups of ports:

- **Link to Monitor (LtM)**—The Link to Monitor group consists of one port or one trunk group. The switch monitors the LtM for link failure.
- **Link to Disable (LtD)**—The Link to Disable group consists of one or more ports and trunk groups. When the switch detects a link failure on the LtM, it automatically disables all ports in the LtD. When the LtM returns to service, the switch automatically enables all ports in the LtD.

Spanning Tree Protocol with UFD

If Spanning Tree Protocol (STP) is enabled on ports in the LtM, then the switch monitors the STP state and the link status on ports in the LtM. The switch automatically disables the ports in the LtD when it detects a link failure or STP BLOCKING state.

When the switch determines that ports in the LtM are in the FORWARDING state, then it automatically enables the ports in the LtD, to fall back to normal operation.

UFD Configuration Guidelines

This section provides important information about configuring UFD.

- Only one Failure Detection pair (one group of Links to Monitor and one group of Links to Disable) is supported on the switch (all VLANs and Spanning Tree Groups).
- An LtM may contain either one port or one Multi-Link trunk group.

- Ports that are already members of a trunk group are not allowed to be assigned to an LtM.
- A port cannot be added to a trunk group if it already belongs to an LtM.
- An LtD can contain one or more ports, and/or one or more trunks.

Use the following command to find out how many times link failure was detected on the LtM, how many times Spanning Tree blocking state was detected on the LtM, and how many times UFD disabled ports in the LtD:

```
ex2500(config)# show ufd counters
```

UFD Configuration Example

Figure 14 shows a basic UFD configuration. In this example, NIC 1 is the primary network adapter; NIC 2 is a non-primary adapter. NIC 1 is connected to port 16, and NIC 2 is connected to port 17. Port 2 is connected to a Layer 2/3 routing switch.

The following procedure pertains to the example shown in Figure 14:

1. Configure Network Adapter Teaming on the servers.

2. Assign the Link to Monitor (LtM) ports.

```
ex2500(config)# ufd fdp ltm port 2
```

3. Assign the Link to Disable (LtD) ports.

```
ex2500(config)# ufd fdp ltd port 16
```

4. Turn on Uplink Failure Detection (UFD).

```
ex2500(config)# ufd fdp enable
ex2500(config)# ufd enable
```

Monitoring UFD

The UFD information menu displays the current status of the LtM and LtD, and their member ports or trunks. For example:

```
ex2500# show ufd
Uplink Failure Detection: Enabled
```

```
LtM status: Down
Member      STG      STG State      Link Status
-----
port 2      1        DISABLED       down
```

* = STP turned off for this port.

```
LtD status: Auto Disabled
Member      Link Status
-----
port 16     disabled
```


Part 2

Appendixes

- “Monitoring Ports with Port Mirroring” on page 81 discusses the main tool for troubleshooting your switch—monitoring ports.

Appendix A

Monitoring Ports with Port Mirroring

This appendix explains port mirroring to help you monitor ports and troubleshoot common problems on the EX2500 switch. The following topics are discussed in this appendix:

- Port Mirroring Overview on page 81
- Configuring Port Mirroring on page 82

Port Mirroring Overview


The port mirroring feature in the EX2500 switch allows you to copy traffic from specified ports and forward it to another port for monitoring or packet analysis. The port that receives the copied traffic is called the *monitor* port. The ports being monitored, and the traffic being copied, are considered to be *mirrored*.

The port mirroring feature can be used as a troubleshooting tool or to enhance the security of your network. You can attach a *sniffer*, or packet analysis device, to the monitor port and examine the mirrored traffic without disrupting traffic on the mirrored ports. As an example, an IDS server can be connected to the monitor port to detect intruders attacking the network.

The EX2500 switch can mirror all types of Layer 2 and Layer 3 traffic. Up to four monitor ports can be configured. Each monitor port can receive mirrored traffic from multiple switch ports, but each specific switch port is permitted to be mirrored to only one monitor port. For each mirrored port, you can also specify whether to mirror only ingress traffic (traffic entering the switch port), only egress traffic (traffic leaving the switch port), or both.

Figure 15 shows an example of port mirroring.

Figure 15: Monitoring Ports



As shown in Figure 15, port 2 is acting as a monitor port, receiving mirrored traffic from three other switch ports: *ingress* traffic from port 4, *egress* traffic from port 7, and *both* ingress and egress traffic from port 10. A sniffer could be attached to port 2 in order to monitor the mirrored traffic on ports 4, 7, and 10.



NOTE: Ingress and egress traffic is duplicated and sent to the monitor port after regular switch processing is complete.

Configuring Port Mirroring

To configure port mirroring for the example shown in Figure 15:

1. Specify the monitoring port, the mirroring port(s), and the port-mirror direction.

```
ex2500(config)# port-mirroring monitor-port 2 mirroring-port 4 in
ex2500(config)# port-mirroring monitor-port 2 mirroring-port 7 out
ex2500(config)# port-mirroring monitor-port 2 mirroring-port 10 both
```

2. Enable port mirroring.

```
ex2500(config)# port-mirroring enable
```

3. View the current configuration.

```
ex2500# show port-mirroring
Port mirroring is enabled
Monitoring Ports Mirrored Ports
1                none
2                (4, in) (7, out) (10, both)
3                none
4                none
5                none
6                none
7                none
8                none
9                none
10               none
...
```


Part 3

Indexes

- Index on page 85

Index

Numerics

802.1p priority for QoS	63
802.1Q VLAN tagging	23

A

Access Control Lists. <i>See</i> ACLs.	
accessing the switch	
management interface	3
RADIUS authentication	11
security	10
TACACS+ authentication	14
using the Web Device Manager	5

ACLs

configuration	56
IP Extended ACLs	54
IP Standard ACLs	53
MAC Extended ACLs	53
overview	52
well-known protocol types	54
administrator account	13
alarms, RMON	68
application ports	55
audience for this manual	xi

B

Bridge Protocol Data Units (BPDUs)	32
broadcast domains	21
broadcast storms	
configuring storm control	59
description	59

C

caution icon	xii
CIST	39
Class of Service queue	64
Common Internal Spanning Tree	39
community VLAN	29
configuration examples	
ACLs	56
IGMP snooping	73
MSTP	40
multiple VLANs with tagging adapters	27
port trunking	45
RSTP	36
UFD	77

configuration rules

Fast Uplink Convergence	42
LACP	49
MSTP	39
port mirroring	44
private VLANs	30
RSTP	36
Spanning Tree Groups (STGs)	33
trunking	44
UFD	76
VLANs	26, 44
COS queue	64

D

default password	13
default service levels	62
DHCP, enabling	4
Differentiated Services Code Point. <i>See</i> DSCP.	
Differentiated Services. <i>See</i> DSCP.	
DiffServ. <i>See</i> DSCP.	
DSCP	
Assured Forwarding (AF)	61
Class Selector (CS)	61
COS queue mapping, viewing	62
Expedited Forwarding (EF)	61
overview	60
Per Hop Behavior	61

E

edge ports	36
end user access control	18
EtherChannel	43, 44
as used with port trunking	44, 46
events, RMON	69
EX2500 documentation	xiii
EXCLUDE mode, IGMPv3	73

F

Failure Detection Pair	76
Fast Uplink Convergence	41
fault tolerance with port trunking	44
filtering criteria	54
filters. <i>See</i> ACLs	
frame tagging. <i>See</i> VLANs, tagging.	

H

help, requesting	xiii
high availability, overview	75
history, RMON	67
HP-OpenView	7

I

IBM Director	7
ICMP	54
icons, notice	xii
IEEE standards	
802.1D	31
802.1p	63
802.1Q	23
802.1w	31
IGMP	
description	71
IGMP snooping	71
IGMPv3 snooping	73
snooping configuration	73
SSM	73
INCLUDE mode, IGMPv3	73
Internet Group Management Protocol. <i>See</i> IGMP.	
IP address with Telnet	5
IP Extended ACLs	54
IP Standard ACLs	53
IP subnets with VLANs	21
ISL Trunking	43
isolated port	29
isolated VLAN	29

J

JTAC, contacting	xiii
------------------------	------

L

LACP	
configuration	49
configuration guidelines	49
description	47
reducing timeout	49
laser warning icon	xii
Link Aggregation Control Protocol. <i>See</i> LACP.	
link aggregation. <i>See</i> port trunking.	
Link to Disable	76
Link to Monitor	76
link type	36
logical segment. <i>See</i> IP subnets.	
LtD	76
LtM	76

M

MAC Extended ACLs	53
management interface, configuring	3
mirroring ports	81

monitoring ports	81
------------------------	----

MSTP

configuration example	40
MSTP region	39
overview	39
multi-links between switches, port trunking	43
multiple spanning tree groups	37
Multiple Spanning Tree Protocol. <i>See</i> MSTP.	

N

network management	7
note icon	xii
notice icons, usage	xii

O

OSPF	54
------------	----

P

passwords	13
Per Hop Behavior	61
Per VLAN Rapid Spanning Tree Plus Protocol. <i>See</i> PVRST+.	
PHB	61
platforms supported	xi
port mirroring	81
configuration rules	44
port trunking	43
before configuring	44
configuration example	45
configuration rules	44
EtherChannel	43
fault tolerance	44
monitoring	46
restrictions	46
portchannels. <i>See</i> port trunking.	
ports	
for services	55
monitoring	81
physical. <i>See</i> switch ports.	
private VLANs	29
promiscuous port	29
protocol types	54
PVID (port VLAN ID)	22
PVRST+	
configuration	38
multiple spanning trees	37
overview	37

Q

QoS	
802.1p priority	63
ACLs	52
COS queuing and scheduling	64
DSCP	60
DSCP mapping, viewing	62
EX2500 QoS model	52
overview	51
QoS default service levels	62
storm control filters	59
Quality of Service. <i>See</i> QoS.	

R

RADIUS	
authentication	11
port 1812 and 1645	55
port 1813	55
SSH	18
Rapid Spanning Tree Protocol. <i>See</i> RSTP.	
Remote Monitoring. <i>See</i> RMON.	
RMON	
alarms	68
events	69
history	67
overview	65
statistics	66
routers and port trunking	43
RSA keys	17
RSTP	
configuration guidelines	36
description	35
edge ports	36
link type	36
port state changes	35

S

security	
port mirroring	81
RADIUS authentication	11
TACACS+ authentication	14
VLANs	21
segmentation. <i>See</i> IP subnets.	
segments. <i>See</i> IP subnets.	
service ports	55
SNMP	7
configuring SNMPv3	8
enabling v1 and v2	7
HP-OpenView	7
IBM Director	7
trap host configuration	9

snooping	
configuration	73
IGMP	71
IGMPv3	73
Source-Specific Multicast	73
Spanning Tree Protocol	
BPDUs	32
configuration rules	44
Fast Uplink Convergence	41
MSTP	39
multiple instances	38
multiple spanning trees	37
overview	31
PVRST+	37
RSTP configuration guidelines	36
RTSP	35
Spanning Tree Group (STG) guidelines	33
with UFD	76
SSH	
configuring	17
RSA host and server keys	17
SSM	73
statistical load distribution	43
statistics, RMON	66
storm control filters	59
support, technical, requesting	xiii
switch ports VLANs membership	22
syntax conventions	xii

T

TACACS+ authentication	14
tagging. <i>See</i> VLANs, tagging.	
TCP	54
technical support, requesting	xiii
technical terms	
Assured Forwarding (AF)	61
Class Selector (CS)	61
community VLAN	29
Expedited Forwarding (EF)	61
isolated port	29
isolated VLAN	29
port VLAN identifier (PVID)	23
promiscuous port	29
tagged frame	23
tagged member	23
untagged frame	23
untagged member	23
VLAN identifier (VID)	23
Telnet	5
text conventions	xii
timeout, LACP, reducing	49
troubleshooting ports	81
trunk hash algorithm	47
trunking, ports. <i>See</i> port trunking.	

U

- UDP 54
- UFD 75
 - configuration 77
 - configuration guidelines 76
 - example 75
 - Failure Detection Pair 76
 - monitoring 77
 - overview 75
- UFD with Spanning Tree Protocol 76
- Uplink Failure Detection. *See* UFD.
- user access control 18
- user account 13

V

Virtual Local Area Networks. *See* VLANs.

VLANs

- broadcast domains 21
- configuration rules 44
- default PVID 22
- example showing multiple VLANs 27
- ID numbers 22
- multiple spanning trees 31
- multiple VLANs 23
- port members 22
- PVID 22
- security 21
- Spanning Tree Protocol 31
- tagging 22–28
- topologies 26

W

- warning icon xii